



# Remote Agent Experiment

<http://rax.arc.nasa.gov>

<i>Doug Bernard</i>	( <a href="mailto:Douglas.E.Bernard@jpl.nasa.gov">Douglas.E.Bernard@jpl.nasa.gov</a> )	JPL	(818) 354-2597
<i>Gregory A. Dorais</i>	( <a href="mailto:gadorais@ptolemy.arc.nasa.gov">gadorais@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-4851
<i>Ed Gamble</i>	( <a href="mailto:Edward.B.Gamble-Jr@jpl.nasa.gov">Edward.B.Gamble-Jr@jpl.nasa.gov</a> )	JPL	(818) 354-5387
<i>Bob Kanefsky</i>	( <a href="mailto:kanef@ptolemy.arc.nasa.gov">kanef@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-3514
<i>James Kurien</i>	( <a href="mailto:kurien@ptolemy.arc.nasa.gov">kurien@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-4745
<i>William Millar</i>	( <a href="mailto:millar@ptolemy.arc.nasa.gov">millar@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-0263
<i>Nicola Muscettola</i>	( <a href="mailto:mus@ptolemy.arc.nasa.gov">mus@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-4744
<i>Pandu Nayak</i>	( <a href="mailto:nayak@ptolemy.arc.nasa.gov">nayak@ptolemy.arc.nasa.gov</a> )	NASA ARC	
<i>Nicolas Rouquette</i>	( <a href="mailto:Nicolas.F.Rouquette@jpl.nasa.gov">Nicolas.F.Rouquette@jpl.nasa.gov</a> )	JPL	(818) 354-9600
<i>Kanna Rajan</i>	( <a href="mailto:kanna@ptolemy.arc.nasa.gov">kanna@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-0573
<i>Ben Smith</i>	( <a href="mailto:smith@aig.jpl.nasa.gov">smith@aig.jpl.nasa.gov</a> )	JPL	(818) 393-5371
<i>Will Taylor</i>	( <a href="mailto:taylor@ptolemy.arc.nasa.gov">taylor@ptolemy.arc.nasa.gov</a> )	NASA ARC	(650) 604-3364
<i>Yu-Wen Tung</i>	( <a href="mailto:yu-wen.tung@jpl.nasa.gov">yu-wen.tung@jpl.nasa.gov</a> )	JPL	(818) 354-2673

## Deep Space 1

TECHNOLOGY VALIDATION SYMPOSIUM

*February 8-9, 2000/Pasadena, CA*





# Table of Contents

Extended Abstract .....	1
Remote Agent Technology .....	1
Risks .....	2
Validation Objectives .....	2
Test Program and Results .....	3
Applicability to future NASA missions .....	3
Fact Sheet .....	4
Introduction .....	5
Remote Agent Technology .....	<b>Error! Bookmark not defined.</b>
Overview .....	5
Validation Objectives .....	7
Performance Envelope .....	7
Technology Details .....	8
Subsystem Interdependencies .....	13
The Remote Agent Experiment .....	14
Ground Test .....	<b>Error! Bookmark not defined.</b>
Flight Test .....	<b>Error! Bookmark not defined.</b>
Comparison of Flight and Ground Test Results .....	<b>Error! Bookmark not defined.</b>
Validation Summary .....	<b>Error! Bookmark not defined.</b>
Scaling Up .....	<b>Error! Bookmark not defined.</b>
Lessons Learned .....	31
Future Applications .....	<b>Error! Bookmark not defined.</b>
Acknowledgements .....	<b>Error! Bookmark not defined.</b>
References .....	35
Appendix A: Telemetry Channels .....	<b>Error! Bookmark not defined.</b>
Appendix B: DS1 Technology Validation Power On Times .....	<b>Error! Bookmark not defined.</b>

# Remote Agent Experiment Validation

## EXTENDED ABSTRACT

Remote Agent is a reusable artificial intelligence (AI) software system that enables goal-based spacecraft commanding and robust fault recovery. The Remote Agent accepts high level goals from the operators and on-board software modules such as the autonomous navigator (e.g. “During the next week take pictures of the following asteroids and thrust 90% of the time”). The RA determines a plan of action that will achieve those goals and carries out that plan by issuing commands to the spacecraft. The RA considers the current state of the spacecraft both in planning how to achieve its goals and in executing that plan. This allows for robust responses to failures and other contingencies. The RA detects and responds to failures in real-time and if necessary generates a new plan for achieve its remaining goals.

The essential characteristics that differentiate Remote Agent from traditional flight software are that it is *model-based* and *goal directed*. In traditional software programs and expert systems, the programmer decides what the result of a program should be and writes down instructions or rules that attempt to achieve those results. The computer simply executes the instructions or fires the rules with no knowledge of what the intended result was or how it is achieving it. The term *model-based* refers to the fact that Remote Agent is given a model, or general description, of the behavior and structure of the spacecraft it is controlling. The term *goal-directed* refers to the fact that the operator tells the Remote Agent what the desired goal is. Remote Agent uses its models and reasoning algorithms to continuously find actions directed toward the goal, even if failures or anomalies occur.

The Remote Agent’s key contributions are as follows:

- **Goal-based commanding.** The RA is commanded with high-level goals as opposed to time-ordered sequences.
- **Low-level commanding.** Spacecraft operators can also write flexible, high level spacecraft scripts which draw upon Remote Agent’s resource management, diagnosis, and recovery facilities to the extent desired. These allow the same level of control as sequences, but are more flexible and can draw on RA capabilities as appropriate.
- **Closes-loops on-board.** The RA considers current conditions in deciding how to achieve the goals.
- **Fail-operational fault responses.** Rather than always “safing” the spacecraft, it first tries to find a way to achieve the remaining goals in spite of the failure.
- The knowledge that guides the RA’s decisions are encoded in **understandable, maintainable declarative models**. The models are spacecraft specific, but the reasoning engines are mission independent.

## Remote Agent Technology

Remote Agent integrates three separate technologies: an on-board planner-scheduler (PS), a robust plan execution system (Exec), and a model-based fault diagnosis and recovery system variously called MIR or MIR (for Mode Identification and Recovery). These component technologies are described briefly below.

This paragraph is too abrupt. One of the goals of this report as I see it is to hand this as a definitive copy of the RA a technology and the process that we undertook to get to where we did. If so then we need to explain the motivations behind the RA. Nothing has been mentioned as to what the RA does and why. References at the end of the report ought to also be here as part of the introduction. A pointer to the RAX architecture ought to be here also, with a clear indication of the difference between RA as a technology and RAX as an instantiation of this technology.

*Planner Scheduler*—The Planner/Scheduler (PS) generates the plans that RA uses to control the spacecraft. Given the initial spacecraft state and goals, PS generates a set of synchronized high-level activities that, once executed, will achieve the goals. PS consists of a heuristic chronological-backtracking search engine operating over a constraint-based temporal database. PS begins with an incomplete plan and expands it into a complete plan by posting additional constraints in the database. These constraints originate from the goals that the Ground can generate and from constraint templates (e.g. the camera must be pointed at an asteroid to take a picture of it) stored in a model of the spacecraft. PS queries domain-specific planning experts (specialized software modules such as Deep Space One’s navigation system) to access information that is not in its model..

*Smart Executive*—Exec is a reactive, goal-achieving, control system that is responsible for:

- Requesting and executing plans from the planner.
- Requesting/Executing failure recoveries from Livingstone
- Executing goals and commands from human operators.
- Managing system resources.
- Configuring system devices.
- System-level fault protection.
- Reach and maintain safe-modes as necessary.

Exec is goal-oriented rather than command-oriented. We define a goal as a state of the system being controlled that must be maintained for a specified length of time. As a simple example, consider the goal: keep device A on from time x to time y. If Exec were to detect that device A is off during that period, it would perform all the commands necessary to turn it back on. Exec controls multiple processes in order to coordinate the simultaneous execution of multiple goals that are often inter-dependent. In order to execute each goal, Exec uses a model-based approach to

create a command procedure, which is often complex, designed to robustly achieve the goal.

Having "Livingstone" as a heading and talking about MIR without clarifying that they're one and the same is confusing.

*Livingstone*—The MIR inference engine provides mode identification (diagnosis) and mode reconfiguration (recovery) functionality. To track the state of each component (called a mode) in the spacecraft MIR eavesdrops on commands that are sent to the spacecraft hardware by Exec. As each command is executed, MIR receives observations from spacecraft's sensors, abstracted by monitors in the spacecraft's control software. MIR combines these commands and observations with declarative models of the spacecraft components to determine the current state of the system and report it to Exec. If failures occur, MIR uses the same model to find a repair or workaround that allows the plan to continue execution.

The key idea underlying model-based diagnosis is that a combination of component modes is a possible description of the current overall state of the spacecraft only if the set of models associated with these modes is consistent with the observed sensor values. This method does not require that all aspects of the spacecraft state are directly observable, providing an elegant solution to the problem of limited observability.

Why are we talking about risk without talking about IN DEPTH about a. testing b. validation c. advantages of the model based approach ? This is plain silly.

## Risks

The Remote Agent poses the same kinds of risks as conventional flight software: if the RA makes an incorrect decision due to a software bug or incorrect knowledge in the domain model it could endanger the spacecraft, or even lose the mission.

The perceived risk for the Remote Agent is often higher than for traditional flight software. The reason most often cited is that the RA decides autonomously how to achieve its high-level goals based on current spacecraft conditions. This is often perceived as risky, since it is harder to predict what the spacecraft will do. However, this behavior is not qualitatively different from what fault protection or attitude control do now. They all command the spacecraft based on current state information. The behavior of the RA can be predicted, within an envelope, just as the behavior of ACS or fault control can be predicted within certain bounds. Confidence in the RA's responses can be obtained through testing, just as confidence in ACS and fault protection is obtained now. The only real difference is quantitative: the RA has a larger range of input conditions and responses.

The RA's risks are mitigated in the same way as the risks of flight software, fault control, and ACS: through good design and thorough testing. However, this is where the scale of the RA poses challenges. Fault protection has a short list of responses that can be tested exhaustively; the RA has millions. ACS is governed by control laws for which there are validation methods; the RA is governed by domain models and reasoning algorithms for which there were no validation methods.

We developed new validation methods that were suitable for the scope of the Remote Agent Experiment. These were effective at finding and eliminating bugs in the domain knowledge, and no such bugs were encountered during the experiment. Two bugs in the reasoning engines were discovered during the experiment. We suggest additional testing methods that would be able to detect such bugs on future missions. Testing experiences, lessons, and suggestions for future missions will be discussed in the main body of the paper.

Above is a great advertisement for NOT using the RA. It needs to be rewritten not to emphasize the negatives of the validation and testing efforts that we undertook (which were very weak) but to highlight what the technology can do for risk mitigation. For example, nothing above talks about domain knowledge validation being easier in a declarative model based approach that the RA was built on compared to the conventional model of writing control software which has embedded domain constraints which are hard to understand and separate from the control algorithm. This does nothing to say that such declarative approaches are actually risk mitigators for long term missions simply for the reasons mentioned above. When a declarative model is available it is anticipated that a new mission operator will be in a much better state to understand the dynamics of the s/c by looking at the models than by seeing lines of code in a more conventional approach.

The above section ought to be revised or scrapped altogether.

## Validation Objectives

The validation objectives for RA were broken down into specific objectives for each of the three engines as follows.

*Planner*—The validation objectives for the planner were to: (a) generate plans on-board; (b) reject low-priority, unachievable goals; (c) replan following a failure; (d) generate back-to-back plans; and (e) enable modification of mission goals from ground.

*EXEC*—The EXEC validation objectives were to: (a) provide a low-level commanding interface; (b) initiate on-board planning; (c) execute plans generated both on-board and on the ground; (d) recognize and respond to plan failure; and (e) maintain required properties in the face of failures.

*Livingstone*—The validation objectives for this subsystem were to: (a) confirm executive command execution; (b) demonstrate model-based failure detection, isolation, and recovery; and (c) demonstrate ability to update MIR state via ground commands.

### *Test Program and Results*

The Remote Agent Experiment (RAX) consisted of using the RA technology to operate the DS1 spacecraft for several days. Recall that the RA consists of general-purpose reasoning engines (EXEC, MIR, Planner), and mission-specific *models* that provide the operations knowledge and constraints that the engines reason about in order to make decisions and command the spacecraft. For purposes of validating the RA technology, it was not necessary to develop models for all of the DS1 subsystems and operating modes. We developed an operations scenario based on DS1 active cruise mode, and developed the minimum set of models needed to perform that scenario under all operating conditions and respond to a selected set of faults. Developing additional models would have incurred additional expense without increasing our ability to validate the RA technology. Above is BOGUS. We didn't develop the models for the entire s/c because a. we were hampered by time and b. we were an experiment. It is premature to say that it was "not necessary" to develop the full scale model set. How can we prove this ?

The scenario goals were to execute an IPS thrust arc, acquire optical navigation images as requested by the autonomous navigator, and respond to several simulated faults. The faults included minor ones that could be responded to without disrupting the current plan, and more serious faults that required generating a new plan to achieve the remaining goals. These faults were as follows:

- MICAS switch stuck-on
- communication failure on 1553 bus (PASM RT)
- ACS thruster valve stuck-closed

The RA successfully demonstrated 100% of its validation objectives during the experiment.

### *Applicability to future NASA missions*

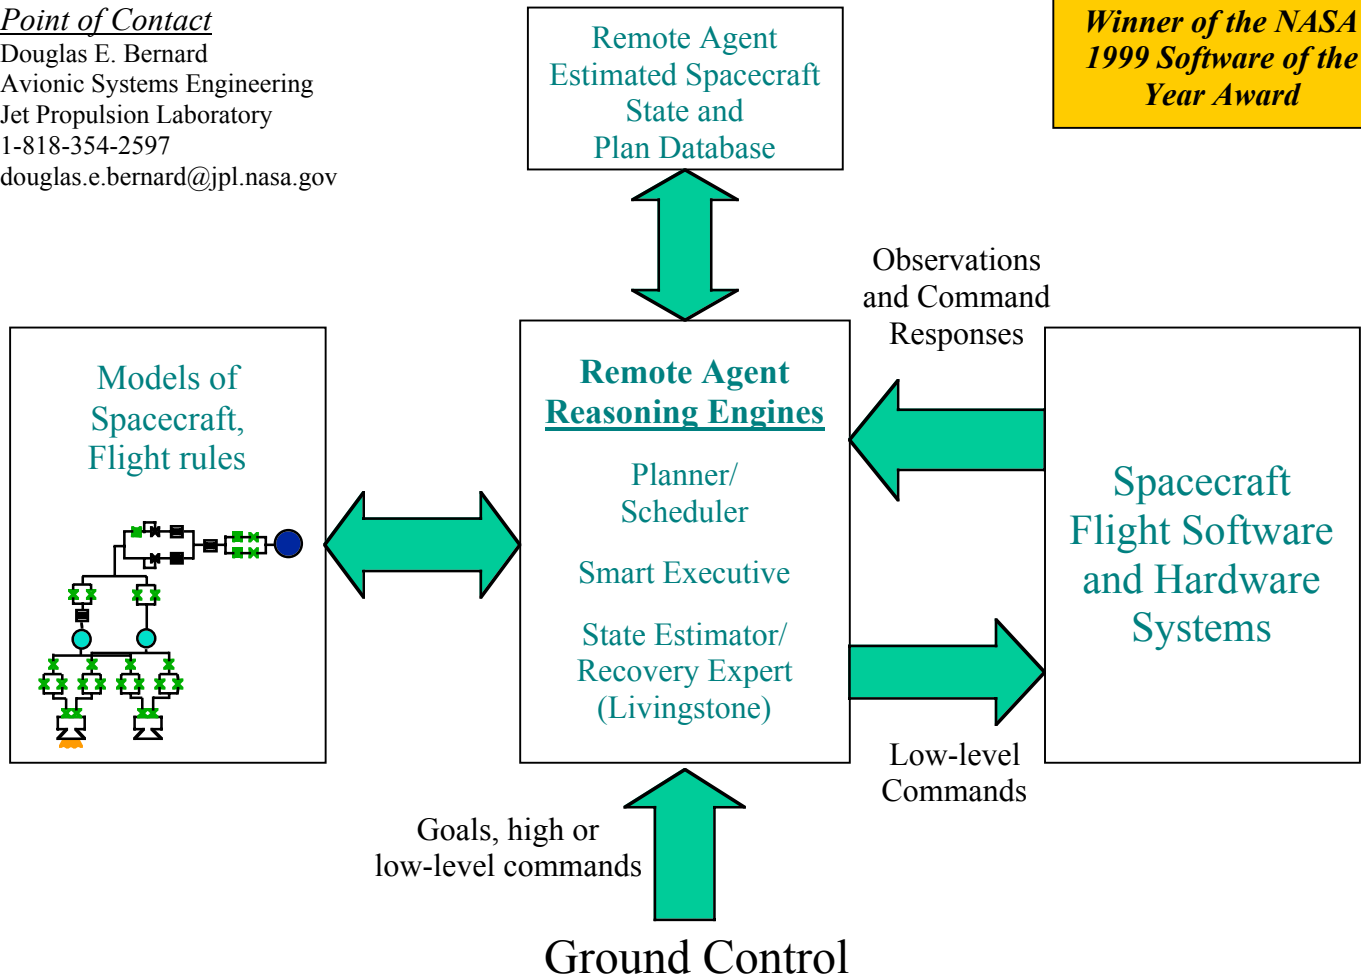
The Remote Agent technology is applicable to any future NASA mission that desires or requires autonomous operations. The RA reasoning engines can be used as-is on future missions. New domain models would be required for each mission.

# Remote Agent Experiment Fact Sheet

## Point of Contact

Douglas E. Bernard  
Avionic Systems Engineering  
Jet Propulsion Laboratory  
1-818-354-2597  
douglas.e.bernard@jpl.nasa.gov

**Winner of the NASA  
1999 Software of the  
Year Award**



## Validation Objectives

- ✓ •Initiate and generate flexible plans on-board
- ✓ •Reject low-priority, unachievable goals
- ✓ •Execute plans generated both on-board and from Ground
- ✓ •Confirm execution of commands
- ✓ •Demonstrate model-based failure detection and recovery
- ✓ •Maintain required spacecraft states in the face of failures
- ✓ •Re-plan following a failure
- ✓ •Generate back-to-back plans
- ✓ •Modify mission goals from Ground
- ✓ •Execute low-level commands from Ground
- ✓ •Update estimated spacecraft state database from Ground

## Capabilities

- Robust Goal-based commanding
  - Planner expands high-level goals into flexible plans
  - Executive decomposes plans into low-level spacecraft commands and monitors that the states commanded to are achieved and maintained
- Fail-operational model-based fault recovery
  - Livingstone identifies faults and suggests recoveries that the Executive uses to continue plan execution
  - If necessary, Executive requests the Planner to generate a new plan in light of failure

## Applicability to future missions

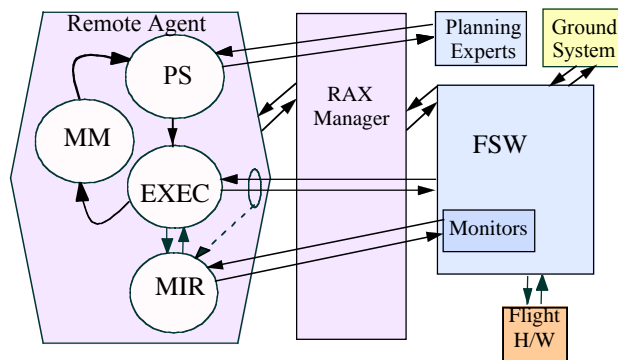
Remote Agent technologies are generally applicable to mission that benefit from highly autonomous operation and are currently being applied to prototypes of future NASA missions including a space-based interferometer and an in-situ propellant production plant.

## INTRODUCTION

Remote Agent is a model-based, reusable artificial intelligence (AI) software system that enables goal-based spacecraft commanding and robust fault recovery. This report describes the Remote Agent technology, its development and test history and the DS1 flight experiment in which Remote Agent was validated. Whenever feasible, this report attempts to give guidance on how Remote Agent can be fruitfully employed in future science missions. We also highlight further technology developments and operational applications of the technologies that we are currently pursuing.

### Technology Overview

The Remote Agent (RA) integrates three separate Artificial Intelligence technologies: automated planning and scheduling, robust multi-threaded execution, and model-based fault diagnosis and recovery.



**Figure 1: Remote Agent architecture**

*Remote Agent Architecture*—The RA architecture and its relation to flight software are shown in Figure 1. Viewed as a black-box, RA issues commands to real-time execution flight software (FSW) to modify spacecraft state, and receives data from the spacecraft through a set of monitors that filter and discretize sensor values. The RA itself is comprised of three main components: a Planner/Scheduler (PS), a Smart Executive (EXEC), and Livingstone, a Mode Identification and Reconfiguration module also known as MIR. An additional component, strictly related with PS, is the Mission Manager (MM). In addition the RA team provided a clean interface to the rest of the FSW. The RAX manager (RAXM) mediated all communication between RA and the flight software and was included from the outset in the design of the DS1 FSW. The RAXM provided a messaging conduit between the RA and the rest of the FSW including interfaces to the Planning Experts, the Monitors and to the Real Time Sequencer. Using such a mechanism allowed the RA to be cleanly bundled on top of the FSW much later in flight and yet also allowed a clear

methodology for testing and validating the software on the ground.

We now describe the main functionalities provided by RA and how each individual RA component participate to the overall picture. We will do so by giving concrete examples of commanding and operations relative to DS1.

As we will discuss later, RA can operate at different levels of adjustable autonomy, allowing ground operators to interact with the spacecraft with immediate commands to the flight software, if needed. However, one of the most unique characteristics of RA, and a main difference with traditional spacecraft commanding, is that ground operators can communicate with RA at the *goal level* rather than having to formulate detailed sequences of timed commands. Goals are stored in the Mission Manager in a *mission profile* covering an extended period of time. In principle, a mission profile could contain all goals for a mission, requiring no further uplink from ground. More realistically, mission operations will want to change goals (e.g., the schedule of communications with DSN can be modified on a week by week basis). This can be easily accomplished by uplinking commands to edit the mission profile. Goals typically contain very little detail of how they should be achieved. For example, for the DS1 Remote Agent Experiment the only goals in the mission profile were “Perform AutoNAV orbit determination (OD) activities for 1 hour every day” and “Thrust the IPS engine for at most 12 hours”.

To translate high-level goals into a stream of commands to flight software, RA follows a two-step process. In the first step, MM selects goals for the next commanding horizon (typically covering several days) and sends them to PS. PS uses its model of the spacecraft to determine which detailed tasks should be selected and scheduled to achieve the goals. For example, in order to perform an OD PS determines from the model that pictures of beacon asteroids need to be taken. In order to select these asteroids, the model instructs PS to interrogate the AutoNAV software as a *planning expert*. In general, PS will rely on several specialized services provided by software modules external to RA. In DS1 both AutoNAV and ACS provided information to PS that was incorporated into plans. Going back to our example, observing an asteroid translates, according to the PS model, into taking a series of images of it with MICAS. Therefore PS schedules a “MICAS take OPNAV images” task. Moreover, the model instructs PS that while images of an asteroid are taken, the attitude of the spacecraft must be compatible with the MICAS camera pointing at it. If this is not the case, the PS model instructs PS to schedule an appropriate turn changing the attitude from the previous one to the desired one.

The brief example above points out another fundamental characteristic of all RA components: their fundamental reliance on explicit, declarative *models* of the spacecraft. Although their level of detail varies between the different components, RA models tend to be fairly abstract and tend to focus on system level interactions rather than on the detail operation of individual subsystems or components.



This has two advantages. Firstly, it provides a method to capture system level knowledge in a form that is directly usable to command a spacecraft, without having to go through the costly and error prone translation of system requirements into flight software. In the traditional approach, system requirements are at best translated into flight rules that can be used to check the validity of command sequences but not to generate them in the first place. Secondly, more abstract models are less susceptible to changes when a detailed understanding of the behavior of each subsystems is gained during spacecraft development. Although they need to be adjusted to the new finding, the abstract models usually remain structurally unchanged, and so remain the synthesis procedures that RA component use to generate command loads.

Going back to our example, once PS has generated a plan for the next commanding horizon, EXEC receives it and incorporates it into the queues of tasks that it is currently executing. Tasks generated by PS tend to be still fairly abstract. EXEC's responsibility is to synchronize the parallel execution of the plan tasks according to the specifications contained in the plan, and to further decompose each task, one at a time, into more detailed steps. This task decomposition eventually results into individual commands being sent, one at a time, to the flight software. For example, the abstract task "MICAS take OPNAV images" is decomposed into commanding MICAS to take a number of snapshots while checking that MICAS is kept "ON" during the entire process.

Besides its goal-directed commanding and model-centered approaches, RA puts particular emphasis on robustness of execution and flexibility of response to faults. The model identification (MI) component of MIR observes EXEC issuing commands, receives sensor observations from monitors, and uses model-based inference to deduce the state of the spacecraft and provide feedback to EXEC. The other component of MIR, mode reconfiguration (MR), serves as a recovery expert, taking as input a set of EXEC constraints to be established or maintained, and recommends a recovery action to EXEC that will achieve those constraints. MIR provides both the MI and MR function using a single core algorithm and a single declarative model.

Fault protection in RA happens at two different levels.

- *Low-level fault protection loop:* This involves EXEC and MIR in the context of executing a single PS-generated task. Suppose that EXEC is commanding MICAS power on in order to ensure that micas is on during the "MICAS take OPNAV images" PS task. It does so by sending an appropriate command to the power driver. MI observes the command and, on the basis of its previous state estimate and its models, predicts the likely next state in which the system will be. This prediction provides a qualitative description of the sensor readings MIR should observe from the spacecraft (e.g. the switch sensor and current sensor should be consistent with MICAS being on). If the expected observations are not received, MI

uses its model to hypothesize the most likely cause of the unexpected observations in terms of failures of the spacecraft's components. The information about the new state of the spacecraft hardware is sent to EXEC, which now asks MIR for an action to correct the problem. MIR now activates MR, which, using the same model, determines the least-cost system state that satisfies EXEC's request and is reachable from the fault mode. MIR then gives EXEC the first action in a possible sequence that will take the system to that state. Such a recovery might involve resetting a device, attempting a command again, or a complex reconfiguration of the spacecraft to enable a functionally redundant system. EXEC executes the recovery action, under the watchful eye of MIR, and receives further actions from MIR if needed by the recovery process. When the recovery is complete, EXEC continues executing the PS task in a nominal fashion. Note that during this entire process the original PS task is still active and in a "nominal" state. This depends on the time allocated to the task including enough slack to tolerate variations during execution that can be handled by low-level fault protection.

- *High-level fault protection loop:* this involves EXEC and PS. Going back to our example, assume that all recovery actions suggested by MR failed and no more recovery actions are available. MIR therefore infers that MICAS is unusable and communicates this to EXEC. This means that there is no way to execute a command necessary for the success of the "MICAS take OPNAV image" task. Moreover, the assumed conditions for other tasks that may be present in the plan in the future may now be invalidated. Therefore EXEC terminates task execution with a failure, discards the rest of the plan and immediately commands the spacecraft to enter an appropriate "RA standby" mode.<sup>1</sup> It then activates PS by communicating to it the current state of the spacecraft and asks for a new plan. After receiving the initial state from EXEC and the goals from MM, PS generates a new plan that achieve the goals as best as possible within the new, degraded configuration of the spacecraft. When the plan is ready, PS sends it to EXEC. EXEC now exists for the "RA standby" state and resumes normal operations by starting the execution of the new plan.

With the above capabilities, RA allows implementation of *fail-operational* behaviors under a much broader range that it is possible in traditional spacecraft commanding. Traditionally only critical sequences (e.g., Saturn Orbit Insertion for Cassini) are designed to tolerate a large number of faults without requiring "safing" of the spacecraft. This depends on the cost of analysis and implementation of these sequences. Therefore in less critical phases of a fault event usually requires the intervention of the ground operations team to correct it. With RA the cost of implementing these

<sup>1</sup> Note that this is a standby situation only from the perspective of RA. From the point of view of FSW, "RA standby" mode is not a fault mode and does not require the intervention of FSW fault protection.

scenarios is significantly reduced, making possible an increase of mission productivity and a reduction of cost of operations.

### Detailed Validation Objectives

Validation of a technology with the complexity and the pervasive systemic impact of RA required attention to several different aspects and dimensions.

The first and most obvious objective was to validate the fact that RA could indeed autonomously command a system as complex as a spacecraft for an extended period of time. This translated into the following list of objectives for each RA component.

#### PS/MM

- generate plans on-board the spacecraft
- reject low-priority unachievable goals
- replan following a simulated failure
- enable modification of mission goals from ground.

#### EXEC

- provide a low-level commanding interface
- initiate on-board planning
- execute plans generated both on-board and on the ground
- recognize and respond to plan failure
- maintain required properties in the face of failures.

#### MIR

- confirm executive command execution
- demonstrate model-based failure detection, isolation, and recovery;
- demonstrate ability to update Livingstone state via ground commands.

Other validation objectives addressed the impact of the introduction of RA into a “traditional” spacecraft software architecture. From the outset RA was designed to work in conjunction with existing FSW modules and not to replace it. As a result the fidelity of control provided by the RA depends on the scope and detail of the models of the spacecraft it has. The challenge was to demonstrate that such cooperative arrangement with FSW could indeed be carried out. This consisted in modeling within RA only a specific set of spacecraft subsystems and allowing conventional techniques of FSW control to deal with the remaining control modes of the craft. While there are no software or architectural limitations which would disallow RA to command all subsystems for an extended period of time, the fielding of RA on DS1 was also meant to provide a credible demonstration of the fact that autonomy concepts could be applied within a well defined scope.

Level	Ground System	On-Board PS	On-Board EXEC
-------	---------------	-------------	---------------

1	Prepare real-time commands	None	None (executed w/o Exec involvement)
2	Prepare sequence	None	Execute sequence
3	Prepare plan, upload to EXEC as script	None	Execute plan; “Scripted mode”
4	Prepare plan, upload to planner as goals	Confirm and pass thru the planner	Execute plan; “Planner Mode”
5	Prepare plan including some unexpanded goals	Complete the plan	Execute plan
6	Define goals	Prepare plan	Execute plan

**Table 1: Autonomy levels of RA**

Even within the scope of the autonomy demonstration, it was important to show that adopting RA was not an “all or nothing” proposition and that RA could be commanded with adjustable levels of autonomous operations. Table 1 shows the possible RA autonomy, all the way from generating low-level sequencer commands using conventional methods which are executed by the EXEC onboard (autonomy level 2), to preparing a plan on the ground and uplinking it to the spacecraft for execution (autonomy level 4), finally to providing closed-loop planning and execution on the spacecraft (autonomy level 6). The DS1 autonomy experiment was designed from the outset to start with level 3 as a confidence building measure and then to migrate to level 6 on providing adequate confidence of the plans generation on the ground and execution onboard.

The final set of validation objectives involved the development process for autonomy software. This involved a number of separate items:

- integration of RA with the DS1 flight software, a large and complex system in itself written in a language (C) different from the RA (Lisp);
- adaptation of RA models and scenarios to reflect operational constraints imposed by the flight team even late in the development process;
- achievement of high-level of confidence by the DS1 spacecraft team by going through a rigorous test regimen dictated by the team on high fidelity testbeds.

We will discuss the level of achievement of achievement of these validation objectives in the section on the Remote Agent Experiment.

### Performance Envelope

Note that these performance and resource figures refer to RA as flown on Deep Space 1 in 1999 in Lisp. Each of the RA engines has been or is being re-architected and ported to C or C++. These new systems may exhibit significantly different performance characteristics.

- **Memory:** 32 Mbytes memory peak, 20 avg.

- **CPU:**
  - RAX ran at priority level just below that of DS1 sequencer (very low).
  - 20% of CPU when planner is idle (only EXEC and MIR are running).
  - 45% of CPU while planner is running (planner, EXEC, and MIR all running).
- **Time to generate plans** depends on plan complexity. RAX plans took 50 to 90 minutes to generate.
- **Telemetry:** 10 bits per second, average.  
This includes notification as each activity in the plan is executed, current diagnosis for each device monitored by MIR, and summary of the planner’s plan generation progress. Similar telemetry would be needed for future science missions.
- **File space:** 140 KB for support files, plus approximately 100 KB per stored plan, depending on plan complexity (proportional to number of activities in the plan). Compressed binary executable was 4MB. At most one plan needs to be stored, though all plans were stored during RAX for validation purposes. RAX also generated a 1MB log file **I disagree**; log files have nothing to do with whether it is/not a science mission. We'll always need log files.

## Technology Details

The Remote Agent consists of general-purpose reasoning engines, and mission-specific domain models. The engines make decisions and command the spacecraft based on the knowledge in the models. This section describes the details of the reasoning engines and how they interact. The DS1 domain models developed for the flight experiment will be discussed in the flight experiment section.

### Planner/Scheduler—

PS provides the core of the high-level commanding capability of RAX. Given an initial, incomplete plan containing the initial spacecraft state and goals, PS generates a set of synchronized high-level activities that, once executed, will achieve the goals. In the spacecraft domain planning and scheduling aspects of the problem need to be tightly integrated. The planner needs to recursively select and schedule appropriate activities to achieve mission goals and any other subgoals generated by these activities. It also needs to synchronize activities and allocate global resources over time (e.g., power and data storage capacity). Subgoals may also be generated due to limited availability of resources over time. For example, it may be preferable to keep scientific instruments on as long as possible (to maximize the amount of science gathered). However limited power availability may force a temporary instrument shut-down when other more mission-critical subsystems need to be functioning. In this case the allocation of power to critical subsystems (the main result of a scheduling step) generates the subgoal “instrument must be off” (which requires the application of a planning step). The PS is able to tune the order in which decisions are made

to the characteristics of the domain by considering the consequences of action planning and resource scheduling simultaneously. This helps keep the search complexity under control. This is a significant difference with respect to classical approaches both in Artificial Intelligence and Operations Research where action planning and resource scheduling are typically addressed in two sequential problem solving stages, often by distinct software systems (see [18]).

Another important distinction between the Remote Agent PS and other classical approaches to planning is that besides activities, the planner also “schedules” the occurrence of states and conditions. Such states and conditions may need to be monitored to ensure that, for example, the spacecraft is vibrationally quiet when high stability pointing is required. These states can also consume resources and have finite durations and, therefore, have very similar characteristics to other activities in the plan. PS explicitly acknowledges this similarity by using a unifying conceptual primitive, the **token**, to represent both actions and states that occur over time intervals of finite extension. More details with examples of the semantics of a token are given further along in this section.

PS consists of a heuristic search engine, the Incremental Refinement Scheduler (IRS) that deals with incomplete or partial plans. Since the plans explicitly represent time in a metric fashion, the planner makes use of a *temporal database*. As with most causal planners, PS begins with an incomplete plan and attempts to expand it into a complete plan by posting additional constraints in the database.

```
(Define_Compatibility
;; computes on SEP_Thrusting
(SEP_Thrusting ?heading ?level ?duration)
:compatibility_spec
(AND
  (equal (DELTA MULTIPLE (Power) (+ 2416
                               Used)))
  (contained_by (Constant_Pointing
                 ?heading))
  (met_by (SEP_Standby))
  (meets (SEP_Standby)))
)

(Define_Compatibility
;; Transitional Pointing
(Transitional_Pointing ?from ?to ?legal)
:parameter_functions
  (?_duration_ <- APE_Slew_Duration (?from
                                     ?to ?_start_time_))
  (?_legal_ <- APE_Slew_Legality (?from ?to
                                  ?_start_time_))
:compatibility_spec
(AND
  (met_by (Constant_Pointing ?from))
  (meets (Constant_Pointing ?to)))
)

(Define_Compatibility
;; Constant Pointing
(Constant_Pointing ?target)
:compatibility_spec
(AND
  (met_by (Transitional_Pointing * ?target
                                  LEGAL))
  (meets (Constant_Pointing ?target *
                              LEGAL)))
)
```

Figure 2: Temporal Constraints in DDL

the mission specific flight rules. This means that (in the case of PS) not only are the core engines (the HSTS Temporal Database (TDB) and the Search Engine) reusable across missions, but that the model that is separate can be manipulated independently of any other piece of the flight code (note that since the heuristics search control information is model dependant, this module would be impacted also). In addition, the richness of the representation and the declarative form of the DDL ensures that mission/systems engineers can have a substantially easier job of understanding and verifying the flight rules in the RA than would have been possible in conventional FSW. These are some of the substantial advantages that the RA brings to a mission.

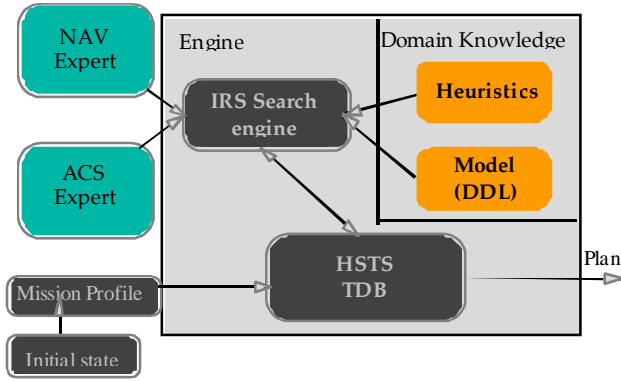
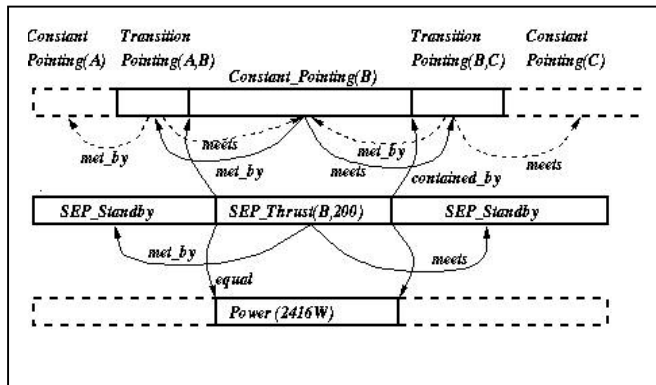


Figure 2: Planner/Scheduler Architecture

Each subsystem in the model is represented in the PS database. Each as a set of dynamic *state variables* whose value is tracked over time. Timelines are treated as instantiations of state variables and are used interchangeably with state variables in this report. Each dynamic state variable can assume one or more values. A token is associated with a value of a state variable occurring over a finite time interval. Each value has one or more associated *compatibilities*, i.e., patterns of constraints between tokens. A legal plan will contain a token of a given value only if all temporal constraints in its compatibilities are satisfied by other tokens in the plan. Fig.3 shows an example of a set of compatibilities with temporal constraints.



The first compatibility indicates that the master token, **Sep\_Thrusting** (when the Solar Electric Propulsion engine is producing thrust), must be immediately preceded and followed by a **SEP\_Standby** token (when the SEP engine is in a standby mode but has not been completely shut off), and it must be temporally contained by a constant pointing token, and the complete thrusting activity requires 2416 Watts of power. The **Constant Pointing** token implies that the spacecraft is in a steady state aiming its camera towards a fixed target in space. **Transitional Pointing** tokens describe an activity when the spacecraft slews. Figure 3 below gives a visual rendering of these compatibilities.

Figure 3: A Plan fragment formed by a DDL model

The timeline approach to modeling is also driven by strong software engineering principals. In a complex domain with different individuals and organisations with varying expertise, timelines provide disparate views of the same domain model across organisational boundaries. For instance the ground team might want to own and access timelines relating to communication coverage and when DSN access is available, while the attitude control team might want to place high level goals on the attitude timeline.

In **Error! Reference source not found.** below we identify four distinct kinds of state variables. A *goal* timeline will contain the sequence of high-level goals that the spacecraft can satisfy (e.g., the Navigate goal described before). Goal timelines can be filled either by ground operators or by on-board planning experts seen by PS as goal generators. For example, in order to generate the portion of the plan that commands the IPS engine, PS interrogates NAV which returns two types of goals: the total accumulated time for the scheduling horizon and the thrusting profile to be followed. These two types of information are laid down on separate goal timelines. Expected device health information over time is tracked by *health* timelines. The expected profile is communicated by EXEC to PS in the initial spacecraft state. EXEC can communicate that the health of a device has changed even if no fault has occurred. Another kind of state variable is an *internal* timeline. These are only used by the planner to internally organize goal dependencies and subgoaling. Finally, an *executable* state variable corresponds to tasks that will be actually tracked and executed by EXEC.

The RAX PS treats all timelines and tokens within a simple, unified search algorithm. This has advantages. The ground team could force certain behaviors of the spacecraft by including in the mission profile explicit tokens on executable timelines. The additional tokens will be treated by PS as goals, will be checked against the internal PS model and missing supporting tasks will be automatically expanded to create a overall consistent plan. This will greatly facilitate the work of the ground team. A plan fragment based on these compatibilities is shown in Fig. 6. For DS1 such models were understandably more comprehensive and complex with more timelines, tokens and compatibilities between differing token types, and required careful consideration during modeling to ensure

interactions between timelines do not result in unanticipated and harmful behaviors generated by the planner.

When a science mission wants to fly the RA planner, the primary set of tasks needed to adapt it to the mission will be:

- perform knowledge acquisition to determine all the spacecraft flight rules.
- encode these flight rules in the DDL model of the spacecraft.
- design the search control heuristics that will be needed to ensure that the planner is able to produce a valid plan within specified resource (time, CPU) bounds.

We do not intend to suggest that models can be or ought to be built in a all or nothing fashion. On the contrary we strongly believe that the process of coming up with a viable plan which encapsulates all the flight rules in the domain is in the incrementality of the process (you "*build some and test some*").

As mentioned previously since the search engine would not need to be adapted but reused, the mission will save costs in revalidating the control system and can confine itself to building and validating the model and search control heuristics. Efforts are underway at NASA ARC to implement automated tools which will ensure that full coverage of the behaviors anticipated by the models (and none other), is simulated during the modeling process. Additional efforts are also underway to automatically generate the heuristics from a given model of the domain. This will further allow mission designers and systems staff to build robust and complex models on their own without relying on the technologists themselves.

Additional details about the planner can be found in [6], [7], [8], [11], [12] and [13].

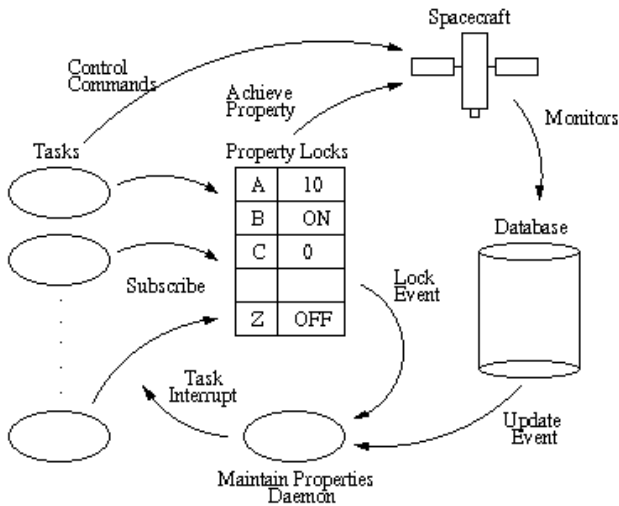
*Executive*—The Smart Executive (EXEC) is a multi-threaded, reactive commanding system. EXEC is responsible for sending the appropriate commands to the various flight systems it is managing. EXEC can replace the traditional spacecraft sequencer or it can be used in conjunction with a traditional sequencer to command a complex subsystem, e.g., interferometer.

EXEC is a multi-threaded process that is capable of asynchronously executing commands in parallel. In addition to the capabilities of a traditional sequencer, EXEC is capable of:

- Simultaneously achieving and maintaining multiple goals, i.e., system states, by monitoring the success of commands it issues and reactively re-achieving states that are lost.
- Conditional sequencing. Commands can be dependent on conditions that occur at execution time.

- Event-driven commanding, as opposed to traditional sequencers that are time-driven. For example, taking a sequence of pictures based on the results of monitoring a range sensor.
- High-level commanding and run-time task expansion. EXEC provides a rich procedural language, Execution Support Language (ESL) [1] in which spacecraft software/model developers define how complex activities are broken up into simpler ones. A procedure can specify multiple alternate methods for goal achievement to increase robustness.
- Sequence recovery. In the event that a command in an executing sequence fails, EXEC suspends execution of the failed sequence and attempts a recovery, either by executing a pre-specified recovery sequence such as reissuing the command or consulting a recovery expert, e.g., MIR. Once the desired state of the failed command is achieved, the suspended sequence is restarted.
- Temporally-flexible sequence (or plan) execution. In order to decrease the probability of a sequence failing, time ranges can be specified for executing and achieving the desired state for each command.
- Resource Management. As a multi-threaded system, EXEC can work on multiple tasks simultaneously. These tasks may compete for system resources within the constraints not already resolved by ground or the planner. EXEC manages abstract resources by monitoring resource availability and usage, allocating resources to tasks when available, making tasks wait until their resources are available, and suspending or aborting tasks if resources become unavailable due to failures (such as a device breaking). See [1] and [2] for a more detailed discussion.

The following figure illustrates key functions of EXEC.



**Figure 6 An Overview of the Remote Agent Executive**

EXEC achieves multiple tasks sending the appropriate control commands (decomposed from high-level commands) to the flight software. The tasks also lock properties that the tasks need to maintain. For example, if task commands a switch ON, the switch property will be locked ON. Monitors (and Livingstone) determine if it is consistent to believe that the switch is ON. EXEC stores this state in its state database. Should the inferred state of the switch change, the database will be updated and an event created signaling a change. If the signaled event violates a property lock, an EXEC property daemon (thread) interrupts those tasks that subscribed to that property lock. EXEC then attempts to achieve the state of the switch being ON using its own recovery mechanism or by consulting a recovery expert, e.g., Livingstone. If the switch cannot be turned ON in time, a hard-deadline that EXEC is tracking is missed, EXEC commands the spacecraft into a safe, wait state while EXEC requests a new plan from the planner that takes into account that the switch cannot be turned ON.

Recoveries may be as simple as sending another command to turn a switch ON, or may be complex, such as when multiple subsystems are tightly coupled. For example, consider two coupled DS1 subsystems: the engine gimbal and the solar panel gimbal. A gimbal enables the engine nozzle to be rotated to point in various directions without changing the spacecraft orientation. A separate gimbal system enables the solar panels to be independently rotated to track the sun. In DS1, both sets of gimbals communicate with the main computer via a common gimbal drive electronics (GDE) board. If either system experiences a communications failure, one way to reset the system is to power\_cycle (turn on and off) the GDE. However, resetting the GDE to fix one system also resets the communication to the other system. In particular, resetting the engine gimbal, to fix an engine problem, causes temporary loss of control of the solar panels. Thus, fixing one problem can cause new problems. To avoid this, the recovery system needs to take into account global constraints from the nominal schedule execution, rather than just making local fixes in an

incremental fashion, and the recovery itself may be a sophisticated plan involving operations on many subsystems.

Domain-code developers use ESL to create high-level commands that EXEC decomposes and executes at run-time depending on the spacecraft state. The following ESL code fragment illustrates multiple methods for achieving IPS thrusting at a desired level depending on the current state of execution. If the IPS is in standby mode, the ACS is commanded to change control modes only after the desired IPS thrust level has been confirmed.

```
(to_achieve (IPS_THRUSTING ips level)
  ((ips_is_in_standby_state_p ips)
   (sequence (achieve (power_on? 'ega-a))
              (command_with_confirmation
                (send-ips-set-thrust-level level))
              (command_with_confirmation
                (send-acb-change-control-mode
                  :acb-tvc-mode))))
  ((ips_in_thrusting_state_p ips)
   (command_with_confirmation
    (send-ips-change-thrust-level level)))
  (t (fail :ips_achieve_thrusting)))
```

EXEC and its commanding language, ESL, are currently implemented using multi-threaded Common LISP. A new version of EXEC is currently under development in C/C++. The internal EXEC code is designed in a modular, layered fashion so that individual modules can be designed and tested independently. Individual generic device knowledge for RAX is implemented based on EXEC's library of device management routines, to support addition of new devices and reuse of the software on future missions.

More details about EXEC can be found in [1], [2], [3] and [8].

*Diagnosis and Repair*— We refer to the diagnosis and repair engine of RA as MIR, for Mode Identification and Reconfiguration. MIR eavesdrops on commands that are sent to the on-board hardware managers by the EXEC. As each command is executed, MIR receives observations from spacecraft sensors, abstracted by monitors in lower-level device managers for the Attitude Control Subsystem (ACS), Bus Controller, and so on. MIR uses an inference engine called Livingstone to combine these commands and observations with declarative models of the spacecraft's components to determine the current state of the system (Mode Identification) and report it to the Exec. Exec may then request that Livingstone return a set of commands that will recover from a failure or move the system to a desired configuration (Mode Reconfiguration). illustrates the data flow between a spacecraft, Exec and Livingstone.

MI is responsible for identifying the current operating or failure mode of each component in the spacecraft, allowing Exec to reason about the state of the spacecraft in terms of component modes, rather than in terms of low-level sensor values. MR is responsible for suggesting reconfiguration actions that move the spacecraft to a configuration that

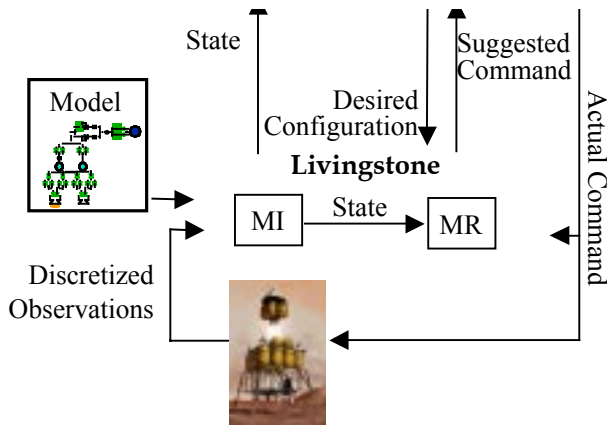


Figure 8 Livingstone Processing Cycle

achieves all current goals as required by PS and EXEC, supporting the run-time generation of novel reconfiguration actions. Though in RA Livingstone is only used to recover following a component failure, its MR capability can be used to derive simple actions to reconfigure the spacecraft at any time. Thus Livingstone can be viewed as a discrete model-based controller in which MI provides the sensing component and MR provides the actuation component. Livingstone uses a single set of models and core algorithms to provide both the MI and MR functions.

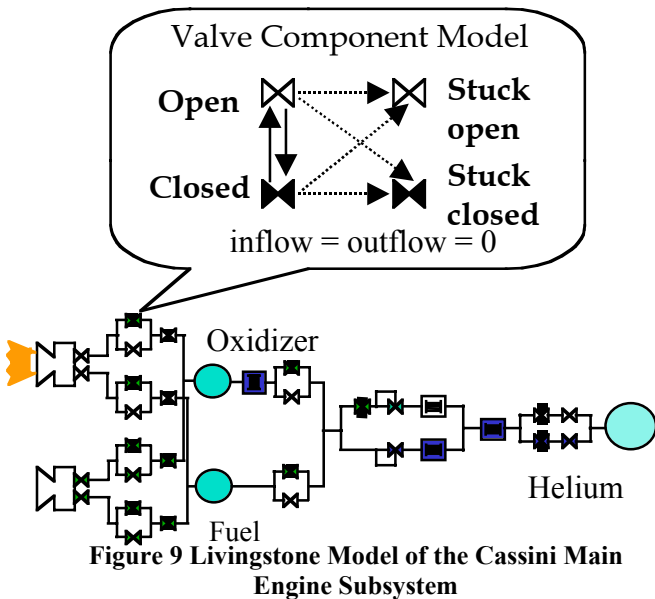


Figure 9 Livingstone Model of the Cassini Main Engine Subsystem

To use Livingstone, one specifies how the components of interest are connected. For each type of component, one then specifies a finite state machine that provides a description of the component's nominal and failure behavior. Figure 5 graphically depicts a Livingstone model of the Cassini main engine subsystem. An important feature is that the behavior of each component state or mode is captured using abstract, or qualitative, models [3, 4]. These models describe qualities of the spacecraft's structure or behavior without the detail needed for precise numerical prediction, making abstract models much easier to acquire and verify than quantitative engineering models. Examples of qualities captured are the power, data and hydraulic connectivity of spacecraft components and the directions in which each thruster provides torque. While such models

cannot quantify how the spacecraft would perform with a failed thruster for example, they can be used to infer which thrusters are failed given only the signs of the errors in spacecraft orientation. Such inferences are robust since small changes in the underlying parameters do not affect the abstract behavior of the spacecraft.

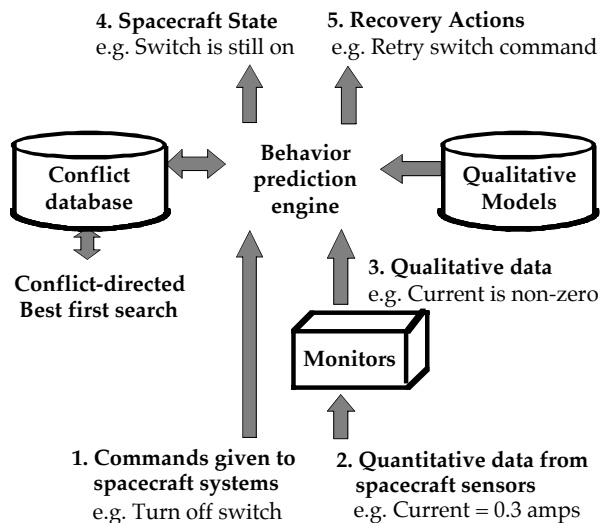
Livingstone's abstract view of the spacecraft is supported by a set of fault protection monitors that classify spacecraft sensor output into discrete ranges (e.g. high, low nominal) or symptoms (e.g. positive X-axis attitude error). One objective of the RA architecture was to make basic monitoring capability inexpensive so that the scope of monitoring could be driven from a system engineering analysis instead of being constrained by software development concerns. To achieve this, monitors are specified as a dataflow schema of feature extraction and symptom detection operators for reliably detecting and discriminating between classes of sensor behavior. The software architecture for sensor monitoring is described using domain-specific software templates from which code is generated. Finally, all symptom detection algorithms are specified as restricted Harel state transition diagrams reusable throughout the spacecraft. The goals of this methodology are to reuse symptom classification algorithms, reduce the occurrence of errors through automation and streamline monitor design and test.

It is important to note that the Livingstone models are not required to be explicit or complete with respect to the actual physical components. Often models do not explicitly represent the cause for a given behavior in terms of a component's physical structure. For example, there are numerous causes for a stuck switch: the driver has failed, excessive current has welded it shut, and so on. If the observable behavior and recovery for all causes of a stuck switch are the same, Livingstone need not closely model the physical structure responsible for these fine distinctions. Models are always incomplete in that they have an explicit unknown failure mode. Any component behavior that is inconsistent with all known nominal and failure modes is consistent with the unknown failure mode. In this way, Livingstone can infer that a component has failed, though the failure was not foreseen or was simply left unmodeled because no recovery is possible. By modeling only to the level of detail required to make relevant distinctions in diagnosis (distinctions that prescribe different recoveries or different operation of the system), we can describe a system with qualitative "common-sense" models which are compact and quite easily written.

Livingstone uses algorithms adapted from model-based diagnosis (see [10]) to provide the above functions. The key idea underlying model-based diagnosis is that a combination of component modes is a possible description of the current state of the spacecraft only if the set of models associated with these modes is consistent with the observed sensor values. Following de Kleer and Williams [9], MI uses a conflict directed best-first search to find the most likely combination of component modes consistent with the



observations. Analogously, MR uses the same search to find the least-cost combination of commands that achieve the desired goals in the next state. Furthermore, both MI and MR use the same system model to perform their function. The combination of a single search algorithm with a single model, and the process of exercising these through multiple uses, contributes significantly to the robustness of the complete system. Note that this methodology is independent of the actual set of available sensors and commands. Furthermore, it does not require that all aspects of the spacecraft state are directly observable, providing an elegant solution to the problem of limited observability. Figure 10 provides a schematic overview of Livingstone's processing.



**Figure 10 Schematic of Livingstone Processing**

The use of model-based diagnosis algorithms immediately provides Livingstone with a number of additional features. First, the search algorithms are sound and complete, providing a guarantee of coverage with respect to the models used. Second, the model building methodology is modular, which simplifies model construction and maintenance, and supports reuse. Third, the algorithms extend smoothly to handling multiple faults and recoveries that involve multiple commands. Fourth, while the algorithms do not require explicit fault models for each component, they can easily exploit available fault models to find likely failures and possible recoveries.

Since RAX, Livingstone has been ported to C++ and significantly improved in the areas of both MI and MR. The improved Livingstone is scheduled to be test flown on both the X-34 and X-37 experimental vehicles. Additional technical details about Livingstone can be found in [4] and at <http://ace.arc.nasa.gov/postdoc/livingstone>

### *Subsystem Interdependencies*

The Remote Agent commands the spacecraft through commands sent to flight software and interacts with a

specific set of subsystems in the FSW. RAX commanded the following systems to demonstrate the validation objectives of the technology

- Ion Propulsion Engine (IPS)
- Miniature Integrated Camera and Spectrometer (MICAS) camera imaging.
- Navigation Expert System (NAV)
- Attitude Planning Expert (APE)
- Attitude Control System (ACS)
- Switches

The RAX Manager (RAXM) is the flight software interface to the Remote Agent experiment (RAX) and isolates the RA software from the rest of the FSW via a set of clean API's. This was the primary use for the RAXM and proved critical in letting the RAX team work in isolation from the DS1 flight team after the redirect.

In addition it provides a terminal in the point-to-point message passing protocol used by the DS1 flight software with all messages passing thru it (see Fig.1) The RAXM in particular is tasked with handling three messages throughout the mission: RAX-START, RAX-STOP and RAX-ABORT; RA software is operational only during the times between a RAX-START and either RAX-STOP or RAX-ABORT. The RAX-STOP is implemented to cleanly terminate the RAX at the end of the experiment under nominal circumstance while the RAX-ABORT is intended to kill the RAX process in the event of an abnormality detected by RAXM. At all other times, the RAXM discards all incoming messages allowing all FSW subsystems that interact with RAX to be ignorant of the RAX state.

When RAX runs, the RAXM handles and dispatches all incoming messages related to RAX -- some of the messages are handled by the RAXM, others are passed through to RAX itself. Similarly, outgoing messages from the RAXM can be due either to the RAXM or to RAX itself.

Like the code for other flight software subsystems, the RAXM is written in the 'C' programming language and is part of the launch load. As a result, the interfaces for RAX needed to be specified early.

The computational resources (CPU fraction, memory space, telemetry buffers and downlink, etc) required by the RAXM when RAX was not running were insignificant. This was by design as a way to mitigate the impact of the RAX technology demonstration on DS1.

### *Preparing Lisp for Flight*

One important aspect of the RAX preparation for flight was the preparation of Lisp for flight. The RAX software development and runtime environment was based on Common Lisp, in particular the Harlequin Lispworks product. The use of Lisp was appropriate given the



background of the RAX developers, the early inheritance of code libraries, and the hardware independence of the high-level software interfaces between RAX and the rest of flight software. However, with the choice of Lisp came some unique challenges. These challenges fell into two rather broad categories: resource constraints and flight software interfaces.

To fit within the 32 MB memory allocation and the CPU fraction constraints, the RAX team thoroughly analyzed their code for memory and performance inefficiencies and employed a "tree-shaking/transduction" process to the Lisp image. The analysis is, of course, common for any high performance software. However, transduction is Lisp-specific and arises from the tight coupling of the Lisp runtime and development environments. Transduction removes the unneeded parts of the development environment, e.g., the compiler, debugger, windowing system. The result is a significantly smaller image, both in terms of file system and runtime memory. During RAX testing, peak memory usage was measured at about 29 MB. Upon completion of the transduction process the RAX Lisp image was compressed by a factor of about 3 to 4.7~MB and uplinked to the spacecraft. On-board decompression was initiated at the start of each RAX run, with the file being inflated directly into the 32~MB RAX memory space. Use of this custom compression drastically reduced the file uplink time and kept the RAX file space usage within the agreed upon limits.

Besides the resource constraints, we also dealt with a complicated flight software interface. The flight software was written in the 'C' programming language and ran on the VxWorks operating system. Lisp and 'C' interacted through Lisp's foreign function interface. This interface was the source of many early problems, primarily caused by discrepancies between data structure alignments assumed by the Lisp and 'C' compilers. These problems were quickly discovered and resolved with the help of an extensive test suite that tested a large number of function parameter variations.

Another problem arose in preparing the Lisp multi-threading system for flight. Originally, the Lisp thread scheduler relied on a high frequency external, periodic wakeup call, issued at interrupt level. However, this went against the design principles of the DS1 flight software. Hence, we had to significantly change Lisp's approach to thread preemption to use a lower frequency wakeup call implemented with flight software timing services.

Most of the late integration problems with RAX Lisp arose because of the VxWorks port. As RAX moved from testbed to testbed, ever closer to the final spacecraft configuration, low-level Lisp problems arose. The problems were consistently of two types: a function assumed by Lisp to be present was not present or a function was present but did not perform as expected by Lisp. The first type of problem was resolved by consistent application of a detailed RAX and

FSW build process. The second type of problem was addressed on a case-by-case basis. Solutions to these problems were made difficult due to the reduced debugging visibility as testbeds assumed the spacecraft configuration. We benefited from the dedicated efforts of both Harlequin and the DS1 FSW team.

## THE REMOTE AGENT EXPERIMENT

During the DS1 mission the Remote Agent technology was validated with an experiment, the Remote Agent eXperiment (RAX). The flight experiment was conducted between May 17<sup>th</sup> 1999 and May 21<sup>st</sup> 1999, and achieved all of the technology validation objectives. However, the story will only be partial without reporting the valuable data gathered during development and testing on the ground. In the case of Remote Agent this is particularly important since the technology is intended as a tool to support system engineering and operations for a mission, rather than simply provide the resulting autonomous capabilities. By quantitatively analyzing the history of development of RAX we can evaluate how well the current state of the technology support its ultimate goals. This can also help identify weak points that require further research and development.

In this section we describe RAX and we attempt to evaluate the development and testing experience with respect to the features of the technology. We first put RAX into the larger perspective of the evolution of the Remote Agent technology. Then we describe the experiment scenarios, the subsystems and fault modes modeled and the expected in-flight behavior. We then discuss how RAX was developed and validated. We support our analysis on the actual record of the problem reports filed in the RAX problem tracking system during development. Finally, we conclude with a description of the flight experiment and an analysis of its results.

*Historical Perspective*—Development of the Remote Agent technology effectively started in May 1995. At that time spacecraft engineers from JPL and Artificial Intelligence (AI) technologists from Ames and JPL started working together on the New Millennium Autonomy Architecture rapid Prototype (NewMAAP), a 6 month effort intended to assess the usability of AI technologies for on-board flight operations of a spacecraft (see [19]). NewMAAP yielded a proof of concept of an autonomous agent that formed the fundamental blueprint for Remote Agent. NewMAAP also helped build the team of technologists that continued development of Remote Agent on DS1.

The successful demonstration of NewMAAP in November 1995 led to the selection of Remote Agent as one of the components of the autonomy flight software for DS1. Between December 1995 and April 1997 the Remote Agent team was part of the DS1 flight software team. This led to the development of the three engines of the RA component technologies and included a substantial speed up of the MIR inference engine (see [4]), the design and implementation of

the ESL language used by EXEC (see [1]), and the design and implementation of the heuristic search engine for PS together with the language to formulate search heuristics. With regard to the overall Remote Agent architecture, we designed and implemented the fault protection protocols, both at the low-level, involving EXEC and MIR and at the high-level, involving EXEC and PS. During this period we acquired much of the high-level system knowledge needed to model DS1 cruise operations (including image acquisitions of beacon asteroids for AutoNAV, timed IPS thrusting and file uplink and downlink) and other DS1 capabilities required for asteroid encounter activities.

In March 1997 the DS1 autonomy flight software was substantially overhauled and DS1 adopted the Mars Pathfinder (MPF) flight software as the basis for its flight software, Remote Agent was re-directed to become an experiment operating for at most six days during the mission on a cruise scenario including AutoNAV orbit determination and IPS timed thrusting. RAX re-used much of the software developed during the previous autonomy flight software phase of DS1. RAX focused on the process of testing each RA component, integrating and testing them into the complete RA, and integrating and testing RA together with the DS1 flight software on the flight processor. Shortcomings found during the development and testing phases required several extensions and re-designs of domain models and the reasoning engines.

In this document we focus solely on RAX and we will make use of the detailed development and testing records maintained during this phase. However, when we will attempt to give our conclusion on technology readiness and costs of the technology, we will do so keeping in mind the entire history of development of Remote Agent. Table 2 below shows the highlights of the RA Experiment starting with the RAX development effort after the redirection of the flight software to MPF. Due to this change, as has been noted elsewhere in this report, a requirement was imposed on us to keep interactions with the DS1 flight team to a minimum. From the beginning, the RAX Manager (RAXM) was identified as being the primary interface to RA and part of the launch load of DS1; delivery of RAXM was initiated by December after negotiating all interfaces with FSW. This was the only significant interaction we had with the DS1 flight team till February 1999. Integration of RAX on the *Radbed* high fidelity testbed was completed during April 1998 which allowed us to understand the timing characteristics of RA in flight. The RAX SDR in September allowed us for the first time to show the DS1 project the progress we were making and explain the expected behavior of RAX during flight. November of the same year barely five months before the experiment, was the first time RAX software ran on a *papabed* after interfacing with the actual FSW; it took another month (and some more) to actually produce a plan and execute it on this testbed. The RAX delivery entered the final deliverable phase in February of 1999 with code development frozen and bug fixes under a strict change control regime. RAX was finally initiated on DS1 on May 17<sup>th</sup> 1999.

Event	Date
Start of RAX development	April 1997
Delivery of RAX Manager to flight software	December 1997
RAX integrated on the flight processor	April 1998
Project Software Delivery Review	September 1998
DS1 launch	October 1998
First run of RAX with FSW on high-fidelity hardware simulation	November 1998
Beginning of M5 DS1 project phase	February 1999
RAX experiment	May 1999

**Table 2: Significant events for the RAX project**

We now describe in detail the DS1 subsystems modeled in RAX and the scenarios on which Remote Agent was exercised during RAX development and testing.

*Domain models*—We only developed domain models for the subsystems and fault modes that were necessary for the experiment. Table XXX describes the timelines modeled by the planner. Tables XXX and XXX list the components and module models developed for MIR. These models captured the following subsystems and resources:

- **Ion Propulsion System**  
Detect and command standby through thrusting states.
- **Attitude Control Subsystem.**  
PS planned attitude changes requested by NAV (IPS attitudes and beacon asteroids) or specified as goals in the mission profile. These attitudes were restricted in the model to slews that maintained the solar panels on-sun. For the experiment, the NAV profiles and goals were specified to further limit the attitudes to either HGA at Earth (the default attitude and the IPS thrust attitude), or MICAS bore-sight at a beacon asteroid.
- **MICAS**  
PS planned data takes and low-voltage power on/off (switch status and commands were modeled, but the switch commands are not actually issued. See the scenario description for why this is so.)
- **Power**  
PS tracked predicted peak-power usage for each activity in the plan (e.g., IPS thrusting, MICAS on), and ensured that the total would never exceed the available power from the solar panels, as predicted by the operations team and supplied in the mission profile

MIR modeled a portion of the power distribution system and its relays in order to confirm operation of the switches commanded by RAX and disambiguate between failures in the power system and erroneous sensor readings. MIR modeled switches not

commanded by RA so that it could request the experiment be aborted if the power system was in a state out of scope for the experiment.

- **Reaction Control System**  
MIR modeled the thruster pallets, thrusters and valves of the RCS system in order to determine the health of the various components from errors in attitude and recommend which control mode to utilize
- **Data System**  
MIR modeled the 1553 bus and a subset of the remote terminal devices on it in order to monitor for remote terminal hangs and recommend resets. Resetting was limited to the PASM instrument. Other remote terminals were modeled in order to allow MIR to request the experiment be aborted if certain out of scope data system problems occurred.
- **Sensors**  
MIR modeled a subset of the switch position and current sensors on-board DS1 as fallible components in order to allow sensor failure as an explanation for unexpected observations
- **Remote Agent**  
PS models aspects of the operation of the Remote Agent itself. For example, the Planner timeline allows PS to plan time for its next planning activity. The Special Activities timeline allows PS to schedule execution of scripts that (unbeknownst to RA) will cause simulated failures on-board the spacecraft.

## *Experiment Scenarios*


The RAX experiment proposal contained of a 12-hour scenario and a 6-day scenario. The 12-hour scenario was designed as a confidence builder for the DS1 project. The 6-day scenario was to be run following successful completion of the 12-hour scenario. Together, the 12-hour and 6-day scenarios demonstrate all RAX validation objectives and they were used for all RAX integration and testing until the beginning of March 1999. In March 1999 the DS1 project levied additional constraints on how the spacecraft could be commanded and specified that RAX should produce 12 hours of thrust or less. We responded by developing a 2-day scenario that met the additional commanding constraints and provided 12 hours rather than 4 days of thrusting. Our ability to quickly develop a new scenario in response to these new constraints was viewed very favorably by the DS1 project. Each of these scenarios is described below.

*Twelve-hour Scenario*— The twelve-hour scenario involves neither on\_board planning nor thrusting with the Ion Propulsion System (IPS). Rather, the plan is generated on the ground, uplinked to the spacecraft, and executed by EXEC and MIR. The scenario includes imaging asteroids with the MICAS camera to support optical navigation, a simulated sensor failure scenario, and demonstration of low-level commanding from a script through RAX to flip a switch. The planning of optical navigation imaging provides the planner the opportunity to reject low-priority, unachievable goals since the optical navigation windows had time only to image a subset of the asteroid goals.

*Six-day Scenario*— The 6-day scenario includes both on\_board planning and operating the IPS, and is a full up test of RA. The scenario is divided into 2 planning horizons. At the start of the scenario, PS generates a plan for the first horizon that included MICAS imaging for optical navigation and IPS thrusting. Execution of the first plan also includes a ground command to modify the goals for the second horizon. At the end of the optical navigation window PS plans to switch off the MICAS camera. However, a stuck on failure injection in the camera switch prevents RA from turning off the camera, leading to a plan failure. Repeated attempts to recover the problem fail. This leads to a replan, which produces a second plan with the camera being left on. The second plan also includes an activity to produce a plan for the second horizon (the third plan in the scenario), which is executed back\_to\_back with the second plan. While the second plan is being executed, the switch failure injection is undone and ground informs MIR that the switch is now fixed. The execution of the third plan includes IPS thrusting, optical navigation imaging, and two simulated failures, a communication failure on the 1553 bus, and a thruster valve stuck closed failure.

The MICAS stuck-on failure demonstrates how MIR and EXEC can make repeated attempts to recover a camera switch until it is deemed permanently stuck. The 1553 bus remote-terminal failure illustrates successful recovery of communication with a device by resetting its remote terminal (RT). In the ACS thruster stuck-closed failure MIR infers from an attitude error and models of the spacecraft dynamics that one of a particular pair of thruster valves is stuck closed. MIR is then able to recommend that no matter which one of the two valves is stuck, switching ACS control modes will mitigate the problem.

*Two-Day Scenario*—In March 1999, the DS1 project analyzed the 6-day plan and decided that RA should not switch the MICAS camera off after each use due to concerns about thermal effects. In addition, RA would be required to produce at most 12 hours of IPS thrusting to




**Figure 1213: 12-hour scenario with no SEP and Ground-based Planning.**

ensure that DS1 would be on track for its asteroid encounter in July 1999.

The 2-day scenario was created that is similar to a compressed 6-day scenario, except that the simulated MICAS switch failure was active for the whole duration of the scenario. This prevented RA from ever switching off the camera. Furthermore, the mission profile was adjusted so that PS would produce plans with only about 12 hours of IPS thrusting. This scenario is similar to the standard DS1 cruise phase, which consists of IPS thrusting punctuated with periodic optical navigation activities. This baseline demonstrated RAX's basic commanding capabilities.

This scenario retains the simulated faults that exercise RAX's robust fault response capabilities. Since we could not depend on failures occurring during the experiment,

failures were simulated by injecting false monitor readings consistent with the failures. While simulations are necessary for demonstration, the RAX is fully responsible for responding to real failures within its limited scope occurring during the experiment. To avoid potential conflicts between RAX and the flight software fault protection mechanism (FP) the RAX response threshold is a little lower than that of FP to allow RAX to detect and respond to faults before FP does. If RAX fails to resolve a fault quickly enough, the FP response would be triggered (since the fault is still active). The FP response is to terminate RAX and resolve the fault.



**Figure 1415: First day of the 2-day scenario**




Figure 1617: Second day of the 2-day scenario

### RAX development

RAX was developed on a number of platforms of decreasing processor speed and increasing level of hardware and software fidelity (see Table 3).

In RAX we adopted a continuous integration development process. New software capabilities were first developed on the *UNIX* platform. Before they could be incorporated in a software build and be appropriately tagged, new features or bug fixes had to run to completion a representative set of scenarios. As time progressed, testbeds of higher and higher fidelity became available. As this happened, the requirements for acceptance of software modifications became more and more demanding since the scenarios had to run on all available platforms.

Besides the speed of the processors another factor affecting productivity was the simulated clock speed. The *UNIX*, *Babybed* and *Radbed* platform made use of a low fidelity simulation developed by the RAX team, which essentially only simulated the message traffic and the delays in receiving responses from flight software. This meant that the simulator was allowed to advance the clock at “warp” speed, simulating in a second several minutes or hours of actual elapsed time. Time warping allowed us to run to completion the full six days scenario in less than an hour, tremendously increasing the productivity during development and testing on such lower fidelity testbeds.

Platform	Fidelity	CPU/OS	Hardware	Availability	Speed
DS1 Spacecraft	Highest	Rad6000 VxWorks	Flight	1 for DS1	1:1
DS1 Testbed	High	Rad6000 VxWorks	Flight spares + DS1 sims	1 for DS1	1:1
Hotbench	High	Rad6000 VxWorks	Flight spares + DS1 sims	1 for DS1	1:1
Papabed	Medium	Rad6000 VxWorks	DS1 sims only	1 for DS1	1:1

Radbed	Low	Rad6000 VxWorks	RAX sims only	1 for RAX	1:1
Babybed	Very Low	PowerPC VxWorks	RAX sims only	2 for RAX	7:1
UNIX	Lowest	SPARC UNIX	RAX sims only	unlimited	35:1

Table 4: Development testbeds for RAX

UNIX	August 1997
Babaybed	February 1998
Radbed	April 1998
Papabed	November 1998
Hotbench	March 1999
DS1 testbed	April 1999
DS1 spacecraft	May 1999

Table 5: Dates of RAX readiness on testbeds

Since the higher fidelity testbeds could not be warped in time because of interfaces to the actual FSW code, it induced us to devise reduced length scenarios that would exercise in a few hours of actual clock time most or all of the functionalities included in the full, multi-day flight scenarios. These shorter scenarios led to exercising RAX under stress conditions complementary to those addressed by the formal test process. As a consequence continuous integration over the course of testing and development led to the discovery and correction of a large quantity of RAX software problems.

### Ground Tests

To qualify RAX to run on board the DS1 spacecraft, RAX underwent a rigorous program of formal tests. The tests covered nominal and off-nominal situations, and exercised each Remote Agent component individually, the integrated RAX product, and RAX together with the flight software, at all levels of fidelity available on the ground testbeds.

Autonomous systems like RA pose testing challenges that go beyond those usually faced by more traditional flight software. In fact, the range of possible behaviors exhibited

by an autonomous system is usually very large. This is consistent with the expectation that the system operate robustly over a large range of possible values of system parameters. An exhaustive verification of all situations, however, would require an unmanageably large number of test cases. To make matters worse, the tests should ideally be run on high-fidelity testbeds, which are heavily oversubscribed, difficult to configure correctly, and cannot run faster than real time. For example, in RAX we could run only 10 tests in four weeks on the DS1 *Hotbench*. To cope with these time and resource limitations, we employed a “baseline testing” approach to reduce the number of tests. Moreover, we exploited as much as possible the lower-fidelity testbeds to validate system behaviors for which there was high confidence that the test results would extend to the higher fidelity situations. We used the high fidelity testbeds mostly in nominal situations and under stress conditions that required RAX to guarantee spacecraft safety.

The baseline scenario was the scenario we expected to execute in flight, initially the six day and twelve hour scenarios and subsequently the two days scenario. We tested a number of nominal and off-nominal variations around these scenarios. These covered variations in spacecraft behavior that we might see during execution and changes to the scenario that might be made prior to execution. Changes included variations to the goals in the mission profile, variations in when faults might occur, and variations in the FSW responses.

The architecture of RA allowed us to run certain tests on lower-fidelity testbeds and be confident that their results would hold on higher-fidelity testbeds. Specifically, the RA commands and monitors the spacecraft through well-defined interfaces with the FSW. Those interfaces were the same on all platforms, as were the range of possible responses. Only the fidelity of the responses improved with platform fidelity. This allowed us to exercise a wide range of nominal and off-nominal behaviors on the *Babybeds* and the *Radbed*, test the most likely off-nominal scenarios on the *Papabed*, and test only the nominal scenarios and certain performance and timing related tests on the *Hotbench* and on the *DS1 Testbed*. A special case was testing of the PS component of RA which made extensive use of the *UNIX* testbeds for its functional testing.

The remainder of this section describes the tests on each of the testbeds, and discusses the effectiveness of our testing approach given the benefit of hindsight.

*UNIX*—The PS team made extensive use of the *UNIX* testbed for PS unit testing throughout the formal testing process. Use of the *UNIX* testbed was critical since PS is a computationally intensive task and could not take advantage of time warping. Both in nominal and fault response situations, PS essentially operated as a batch process with practically no reliance on the underlying real-time system (e.g., timing services). This allowed us to repeatedly run a batch of 200 tests with several variations of initial states, goals of the planner and model parameters (e.g., possible turn durations). The tests were repeated for each release of

the RA software, providing a certain measure of regression testing for the PS software.

*Babybed and Radbed testing*—Each of the RA modules devised a test suite of nominal and off-nominal scenarios that isolated and exercised key behaviors in each module. For PS this involved a batch of about 70 tests comprising some of the tests in the *UNIX* batch plus tests devised to test system level responses of PS (e.g., response to invalid initial state files or to an asynchronous kill message sent by EXEC). The repetition of the tests from *UNIX* both validated the complete functional equivalence of PS between *UNIX* and PPC and verified the acceptability of PS performance on the real-time architecture. MIR was exercised on a batch of 110 tests covering the likeliest failure contexts. The PS and MIR tests were used for testing EXEC, and a suite of twenty additional scenarios exercised the system-level interaction of all modules. These tests were run rapidly on the *babybeds* and *Radbed* with time warping. Running a scenario was a time-consuming and error-prone process. To alleviate this, we designed an automated testing tool that accepted an encoded scenario description as input, controlled the simulator and ground tools to execute the scenario, stopped the test when appropriate by monitoring the telemetry stream, and stored all logs and downlinked files for later examination. This rapid data collection led to a total running time of about one week for all tests, since tests could be scheduled overnight and required no monitoring. Analyzing the results of the tests, however, was still a time consuming process. These tests were run after each major RAX software release.

*Papabed testing*—*Papabed* was extensively used during development in order to integrate RAX with the DS1 flight software. In the context of the formal testing process, *Papabed* was used only to run six off-nominal system test scenarios on the “frozen” version of the RAX delivered to flight software for the flight experiment. These off nominal scenarios corresponded to the situations that were most likely or had the potential for highest impact on the outcome of the experiment. No bugs were detected in these scenarios, probably because RA responses to off-nominal situations were well tested on the *babybed*.

*Hotbench and DS1 Testbed testing*—The *Hotbench* and *DS1 Testbed* were reserved for testing the nominal scenarios, and for testing a handful of requirements for spacecraft health and safety. RAX was designed with a “safety net” that allowed it to be completely disabled with a single command sent either by the ground or by on-board FSW fault protection. Hence, the only ways in which RAX could affect spacecraft health and safety was by consuming excessive resources (memory, downlink bandwidth, and CPU) or by issuing improper commands. We tested the resource consumption cases by causing RAX to execute a Lisp script that consumed those resources. We guarded against improper commands by having subsystem engineers review the execution traces of the nominal scenarios, and doing automated flight rule checking. The nominal

scenarios were run in conditions that were as close to flight-like as possible.

*Software change control*—As the date of the flight experiment drew closer, our perspective on testing changed. Throughout 1998 the main goal of testing was to discover bugs in order to fix them in the code. Starting in January 1999 the discovery of a bug did not automatically imply a code change to fix it. Instead, every new problem was reported to a Change Control Board (CCB) composed by senior RAX project members. Every bug and proposed fix was presented in detail, including the specific lines of code that needed to change. After carefully weighing the pros and cons of making the change, the board voted on whether or not to allow the fix. Closer to flight, DS1 instituted its own CCB to review RAX changes.

As time progressed, the CCB became increasingly conservative and the bias against code modifications significantly increased. This is demonstrated by the following figures. In total, 66 change requests were submitted to the RAX CCB. Of these, 18 were rejected amounting to a 27% rejection rate. The rejection rate steadily increased as time passed: 8 of the last 20 and 6 of the last 10 submitted changes were rejected.

The reason for this increase in conservatism is easily explained. Every bug fix modifies a system that has already gone through several rounds of testing. To ensure that the bug fix has no unexpected repercussions, the modified system would need to undergo thorough testing. This is time consuming, especially on the higher fidelity testbeds, so that full revalidation became increasingly infeasible as we approached flight. Therefore, the CCB faced a clear choice between flying a modified RAX with little empirical evidence of its overall soundness or flying the unmodified code and trying to prevent the bug from being exercised in flight by appropriately restricting the scenario and other input parameters. Often, the answer was to forego the change.

*Summary of testing resources*—We conducted about 200 functional tests for PS on *UNIX* (repeated for 6 software releases), about 300 *Babybed* tests (repeated for 6 software releases), 10 *Papabed* tests (run once), 10 *Hotbench* tests (repeated for two releases), and 2 *DS1 Testbed* tests (on the final release) over a period of 6 months with four half-time engineers. This figure includes design, execution, and analysis of the test cases, and development of testing tools.

<u>Subsystem</u>	<u>State Variables</u>	<u>Value Types</u>	<u>Compatibilities</u>	<u>Comments</u>
MICAS	Executable: 2 Health: 1	7	14	Models the health, mode and activity of the MICAS imaging camera. RAX demonstrates fault injection and recovery for this device as part of the 6 day scenario.
Navigation	Goal: 1 Executable: 1 Internal: 1	5	6	To schedule Orbit determination (OD) based on picture taking activity.
Propulsion & Thrust	Goal: 2 Executable: 1 Internal: 1	9	12	Based on thrust schedule generated by the NAV module, the planner generates plans to precisely activate the IPS in specific intervals based on constraints in the domain model and is the most complex set of timelines and subsystem controlled by the planner
Attitude	Executable: 1 Health: 1	4	4	Enables the planner to schedule slews between constant pointing attitudes when the spacecraft maintains its panels towards the sun. The targets of the constant pointing attitudes are imaging targets, Earth (for communication) and thrust direction ( for IPS thrusting.)
Power Management	Goal: 1 Internal: 1	2	1	Allows the planner to ensure that adequate power is available when scheduling numerous activities simultaneously.
Executive	Goal: 1 Executable: 1	2	7	Allows modeling of low level sequences bypassing planner models giving Mission Ops the ability to run in sequencing mode with the RA.
Planner	Executable: 1	2	2	To schedule when the Executive can request the plan for the next horizon.
Mission	Goal: 1	2	2	Allows the Mission Manager and the planner to coordinate activities based on a series of scheduling horizons updatable by Mission Ops for the entire mission

**Table 8: Summary of Planner models for RAX**

<u>MODULE</u>	<u>TIMELINE</u>	<u>TOKEN</u>	<u>DESCRIPTION</u>
ACS	Spacecraft Attitude	constant_pointing_on_sun	Point vector at Target, Solar Panels at Sun
		transitional_pointing_on_sun	Turn vector to Target, Solar Panels at Sun.
		poke_primary_inertial_vector	Small attitude change.
	RCS_Health	rcs_available	Maintain information on thruster status.
	RCS_OK	maintain_rcs	Set and maintain desired RCS mode.
MICAS (Camera)	MICAS_Actions	micas_take_op_nav_image	Take a set of navigation pictures.
	MICAS_Mode	micas_off	Keep MICAS off.
		micas_ready	Keep MICAS on.
		micas_turning_on	Turn MICAS off.
		micas_turning_off	Turn MICAS on.
	MICAS_Health	micas_availability	Ensure MICAS is available for use.
Op-Nav	Obs_Window	obs_window_op_nav	Wait for a specified duration.
	Nav_Processing	nav_plan_prep	Send message to prepare navigation plan.



PASM	PASM Available	pasm_monitor	Monitor the PASM switch.
SEP	SEP	sep_standby	Achieve and maintain IPS standby state.
		sep_starting_up	Achieve and maintain IPS start-up.
		sep_thrusting	Maintain a thrust level.
		sep_shutting_down	Stop thrusting and go to standby state.
	SEP_Time Accum	accumulated_thrust_time	Monitor thrust time accumulated.
	SEP_Schedule	thrust_segment	Specifies desired thrust level and vector.
	SEP_Thrust Timer	max_thrust_time	Set a timer and stop thrusting if time reached.
		thrust_timer_idle	Thrust timer is off.
Planner	Planner_Processing	planner_plan_next_horizon	Request and get next plan from planner.
		script_next_horizon	Run the next scripted plan.
General	EXEC_Activity	exec_activity	Execute a low-level sequence file passed as a parameter.
	EXEC_Eval	exec_eval_watcher	Process a specified script.

**Table 9 Timelines and Their respective tokens by Module (EXEC's perspective)**

<u>Module</u>	<u># in Model</u>	<u>Subcomponents</u>
power relay	12	1 switch, 1 switch sensor
power distribution unit	1	12 relays, 3 power buses, 3 current sensors, 1 remote terminal
generic RT subsystem	3	1 remote terminal (Models RT for devices MIR does not otherwise model)
IPS system	1	1 IPS, 1 remote terminal
thruster pallet	4	2 thrusters (X facing and Z facing)
reaction control system	1	4 thruster pallets
PASM subsystem	1	1 remote terminal

**Table 10: DS1 hardware modeled as modules in MIR**

<i>Component Class</i>	<i># in Model</i>	<i>Modes</i>
ion propulsion system (IPS)	1	Standby, Startup, Steady State Thrusting, Shutdown, Beam Out, Controller Hung, Unknown
remote terminal	6	Nominal, Resettable Failure, Power-cyclable Failure, Unknown
attitude control	1	TVC, X for Y, Z for Y, X for Y Degraded, Z for Y Degraded, X for Y Failed, Z for Y Failed, TVC Failed, Unknown
switch	12	On, Off, Popped On, Popped Off, Stuck On, Stuck Off, Unknown
switch sensor	12	Nominal, Stuck On, Stuck Off, Unknown
current sensor	3	Nominal (reported value = real value), Unknown (values unconstrained)
thruster valve	8	Nominal, Stuck Closed, Unknown
thruster	8	Nominal, Unknown
propellant tank	1	Non-empty, Unknown (thruster hydrazine out or otherwise unavailable)
bus controller	1	Nominal, Unknown
vehicle dynamics	1	Nominal (This is a qualitative description of force and torque.)
power bus	3	Nominal (Failure considered too fatal and remote to involve in diagnosis.)

**Table 11: DS1 hardware modeled as components in MIR**

## Ground Tools

To provide adequate coverage and visibility into the workings of the RA onboard, we designed a suite of ground tools interfaced with the real-time telemetry generated by the RA.

The two major goals of the RA ground tools were:

1. To present a summary of the spacecraft status understood easily by the mission operations team.
2. To present enough information about the inner workings of the RA software for the experiment team to quickly recognize and debug problems.

To support these goals, telemetry specific to Remote Agent was downlinked during the test. The RA-specific telemetry included contained:

- Planning events (e.g., planning started, finished, and progress messages)
- Sequence execution events (e.g., plan *p* is starting execution, or plan step *x* started executing at time *t*<sub>1</sub>)
- Mode interpretation events (e.g., valve *v* changed state from nominal to possibly-stuck-closed)
- Messages between RA components and between RA and other DS1 flight software.

This telemetry was event oriented and largely incremental, i.e., interpretation of one message depended upon receipt of earlier messages. This design created some problems as we will detail later. In addition, there were heartbeat (“I’m still alive”) messages from RA components visible in the normal DS1 telemetry. The next sections describe specific ground tools constructed to examine this telemetry.

*Packet View*—Remote Agent telemetry is event-based.<sup>2</sup> PacketView displays these events in a simple one line per message format that is easily understandable to all members of the team. As such, PacketView was the most fundamental of the RA ground tools. The PacketView interface (see Fig.8) displayed four types of telemetry packet for each of the three RA modules and IPCO\_RAXO - on-board messages sent between the three RA components. These were displayed as color-coded scrolling lines of text. Missing packets, detected by gaps in the sequential packet numbering, were displayed as a single line in a unique color. Search capabilities were available based on message type and content.

User selectable dialogs presented “pretty-printed” versions of the single-line packet entries. The “time bar” displayed the most recent “spacecraft sent” Greenwich Mean Time (GMT), the most recent “ground received” GMT time, and a running time since last packet was received.

*ExecView*—ExecView visualized the execution status of plans on board the spacecraft (see Fig. 8). Different kinds of activity associated with different parts of the spacecraft appeared on separate timelines. For each kind of activity, specific planned events were represented as tokens appearing along the timeline. Tokens were color-coded to represent their execution status, i.e., in-the-future, currently-executing, completed and completion-overdue.

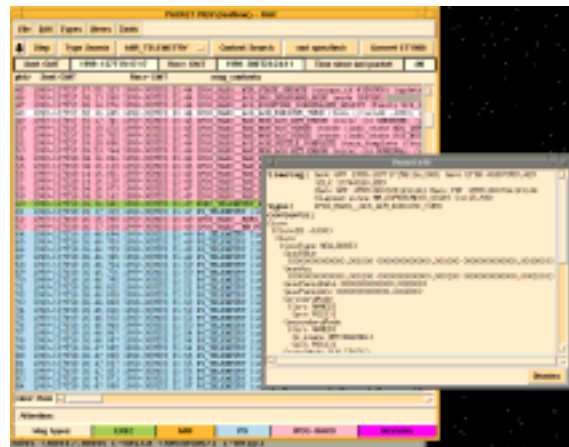


Figure 18: Packetview (telemetry packet) display

As the plan was being executed by Exec onboard the spacecraft, the start and finish times of the activities would be expected to change. Through the constraints, these changes would impact later activities. ExecView would propagate these changes downstream in the schedule, using the same propagation techniques used by the Planner.




Figure 19: ExecView (plan execution status)

ExecView was designed initially as a debugging tool for validating Exec development. As a result it did not have support for handling missing telemetry packets during flight. It produced as a result, some erroneous conclusions during RAX concerning the state of plan execution. To make ExecView more useful, it will have to handle such missing data.

*The Ground Planner*—Of the three technology modules flown as part of the RA, the spacecraft team was least comfortable with PS. To allow the DS1 team to gain confidence in the onboard planner, the RAX team used a ground twin of the planner. The ground planner was identical to the one onboard and was able to duplicate the onboard twin by tapping into the real time telemetry available. It had access to other flight software resources via connection to the papabed. This testbed accurately replicated the software onboard DS1 although it did not replicate the hardware. Of particular importance to the planner were navigation module and beacon asteroid file describing targets for optical navigation and the portion of the attitude control software that predicted the time required to change spacecraft orientation.

The ground planner was a useful tool in predicting the performance of the planner onboard and was especially useful as a confidence builder for mission staff unfamiliar with the working of an autonomous planning agent.

#### *PS Graph*

PS Graph (see Fig. 11) displayed the problem solving trajectory taken by PS for each of the plans generated by the on-board Planner. This took the form of an x-y graph representing the search depth vs. number of search nodes visited for each successive step of the planner's search. The purpose of these plots was to provide a quick summary of how PS problem solving process. For example, a trajectory that visits the same depth level several times while the search node number increases indicates that the planner is backtracking. The persistence of this situation for a large number of steps is an indication that PS may be thrashing and that it will be unlikely to return a solution within the allotted amount of time. Another useful use of the PS Graph plots is the comparison of telemetry data trajectories generated during simulation runs of the ground planner twin.




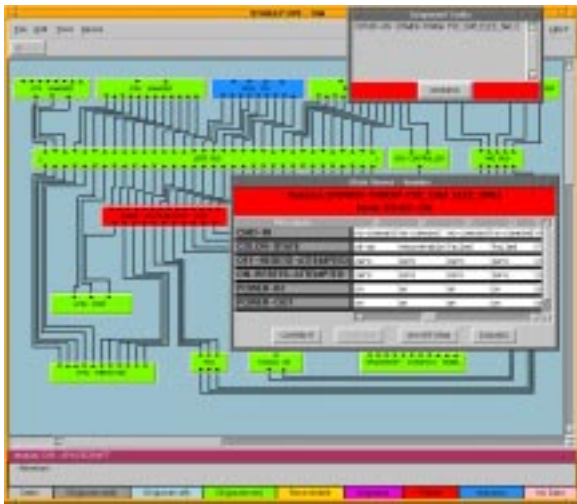
Figure 24: PS Graph (planner progress) display

Although very simple, we were surprised by the power of summarization of this tool and the level of insight that it could provide both during development of RA and during operations in a situation of stress. As we will discuss in the flight experiment section, PSGraph allowed us to monitor an unexpected situation with PS and to quickly identify the likely cause of the problem. In the future, it will be advisable to design several simple visualizations like PS Graph for the reduced ground team supporting an autonomy mission.

#### *Stanley and MIR*

A version of MIR was also run on the ground. The purpose of this was to infer MIR's full internal representation of the spacecraft state from the telemetry which contained a much smaller subset. Specifically, it contained the set of independent variables in MIR's spacecraft model. The Stanley ground tool displayed a hierarchical schematic of the spacecraft's on-board components whose status was driven by the ground MIR (see Fig. 7).

Components could be opened to show more detail or closed. The states displayed were blue (ok - powered off), green (ok - powered on), yellow (recoverable failure), purple (degraded failure), and red (permanent failure). Since Stanley assigned colors to all states, nominal as well as off nominal, it allowed the user to tell at a glance the conditions of the devices. Stanley did not address the issue of displaying continuous values, such as a battery state of charge.



**Figure 25: Stanley (hardware status) display**

In addition to the color changes, detected component faults were reported by popping up an alert box. The alert box allowed the user to click on an entry, resulting in the schematic being opened down to the appropriate hierarchical level to show the local context of the fault. Histories of all state changes, important or not, were available at any time by clicking on components.

*Predicted Events*—In flying an autonomous agent, like the RA, ground operators observing the spacecraft state via its telemetry, may not be in a position to know when certain events are to take place precisely. It was nevertheless important to have a prediction of when the RA planned to take various actions, so that the appropriate subsystem stations at the mission operations center could be staffed for observability. We therefore, generated a Predicted Events File (PEF) file which reported both the low-level commands RA would issue together with the high-level actions RA was asked to take, during plan execution onboard the spacecraft.

#### *Public Outreach via the Web—*

Emailed summaries of events onboard presented in simple English and a Java applet timeline display on the web, patterned after ExecView, were two additional tools to present RA's progress to the public. These tools are interesting because they required an even higher target for simplicity and understandability than did the flight controllers's tools.

Several recent missions have used pagers and email to deliver notifications to the mission operations team. The DS1 ground system, for instance, alerted operators by pager when a given measurement strayed outside a preset range or when fault protection telemetry went into an unusual state. We took this a step further in RAX by producing descriptions of important events in common English. See Figure 12 for an example. The summarized descriptions were automatically posted to the RAX web site (<http://rax.arc.nasa.gov>) and emailed in batches to a

public mailing list. 2000 subscribers received this email during RAX. Terse descriptions were also sent to team members' alphanumeric pagers via email.

We also provided an alternative description of Remote Agent activity (Figure 14) using horizontal timelines patterned after ExecView. This was implemented as a java applet. The timelines in the top window represented major kinds of activity (e.g., attitude or camera-related activity). Along the timelines were tokens indicating particular activities (e.g., a turn) in effect reproducing the plans generated onboard on a user's web browser. Also included were controls to step through the timelines and an event-based summary similar to that provided in email. The most interesting feature of this applet was its ability to show what RA planned to do at any point in the experiment by selecting the event that occurred at that time. This is interesting because the plan changed several times due to simulated faults. Thus it provided an historical overview of RA's re-planning activity and recreated for the general public conditions onboard the spacecraft.

Due to time pressure, the outreach tools were designed to handle the nominal scenario only (including the simulated faults). They did not accurately reflect the RAX software problems that occurred. They did, however, summarize activity during the new scenario without modification. These summaries are still available at the RAX web site at <http://rax.arc.nasa.gov>.



**Figure 26: Timeline applet**

Additional details on the RAX Ground Tools can be found in [14].

### *Flight Test*

RAX was scheduled to be performed on DS1 during a three week period starting May 10, 1999. This period included time to retry the experiment in case of unexpected contingencies. On May 6, 1999, DS1

encountered an anomaly that led to spacecraft safing. Complete recovery from this anomaly took about a week of work by the DS1 team, both delaying the start of RAX as well as taking time away from their preparation for the asteroid encounter in July, 1999. In order not to jeopardize the encounter, the DS1 project also decided to reclaim the third RAX week for encounter preparation, leaving only the week of May 17th, 1999, for RAX. However, to maximize the time to try the more important 2-day experiment, they agreed to go ahead with the 2-day experiment without first doing the confidence building 12-hour experiment. This decision was strong evidence that the DS1 project had already developed significant confidence in RAX during pre\_flight testing.

*Flight Test Part 1*— The flight experiment started on Monday, May 17th, 1999. At 11:04 am PDT, we received a telemetry packet that confirmed that the two days RAX scenario had started on DS1. Shortly thereafter, PS started generating the first plan. The first plan was generated correctly, but not before an unexpected circumstance created some apprehension in us.

Fig. 7 graphically describes the situation with the output of our Plangraph ground tool. The blue trajectory relates to a *Papabed* test that we ran on May 16<sup>th</sup>, 1999 under identical condition to those of the flight test. The green trajectory describes what happened during flight. The deviation in the green trajectory from the 45 degrees, diagonal trajectory means that PS in flight backtracked significantly more than on *Papabed*. Since the conditions on the spacecraft were believed to be practically identical to those on the ground testbeds, there was no apparent reason of for this discrepancy. Subsequently, the cause for this discrepancy was traced back to the spacecraft and *Papabed* differing on the contents of the AutoNAV file containing asteroid goals. Therefore, in flight PS was actually solving a slightly different problem than it had solved on the ground! Thus, this unexpected circumstance allowed us to demonstrate that PS problem solving was robust to last minute changes in the planning goals, increasing the credibility of the autonomy demonstration.

The two days scenario continued smoothly and uneventfully with the simulated MICAS switch failure, the resulting replan, long turns to point the camera at target asteroids, optical navigation imaging during which no communication with DS1 was possible, and the start of IPS thrusting.

However, around 7:00 am on Tuesday, May 18, 1999, it became apparent that RAX had not commanded termination of IPS thrusting as expected. Although plan execution appeared to be blocked, telemetry indicated that RAX was otherwise healthy. The spacecraft too was healthy and in no apparent danger. The decision was made to use EXEC's ability to handle low-level commands to obtain more information regarding the problem. Once enough information had been gathered, the

decision was made to stop the experiment. By this time an estimated 70% of the RAX validation objectives had already been achieved.

*Troubleshooting and Recovery*—By late Tuesday afternoon the cause of the problem was identified as a missing critical section in the plan execution code. This created a race condition between two EXEC threads. If the wrong thread won this race, a deadlock condition would occur in which each thread was waiting for an event from the other. This is exactly what happened in flight, though it had not occurred even once in thousands of previous races on the various ground platforms. The occurrence of this problem at the worst possible time provides strong impetus for research on formal verification of flight critical systems. Once the problem was identified, a patch was quickly generated for possible uplink.

Following the discovery of the problem, we generated a 6-hour RAX scenario to demonstrate the remaining 30% of the RAX validation objectives. This scenario included IPS thrusting, three failure scenarios, and back-to-back planning. This new scenario was designed, implemented, and tested, together with the patch, on *papabed* overnight within about 10 hours. This rapid turn around allowed us to propose a new experiment at the DS1 project meeting on Wednesday. The DS1 project decided to proceed with the new scenario. However, they decided not to uplink the patch, citing insufficient testing to build adequate confidence. In addition, based on the experience on various ground testbeds, the likelihood of the problem recurring during the 6 hour test was deemed to be very low. Nonetheless, we developed and tested a contingency procedure that would enable us to achieve most of our validation objectives even if the problem were to recur.

The DS1 project's decision not to uplink the patch is not surprising. What was remarkable was their ready acceptance of the new RAX scenario. This is yet more evidence that the DS1 project had developed a high level of confidence in RA and its ability to run new mission scenarios in response to changed circumstances. Hence, although caused by an unfortunate circumstance, this rapid mission redesign provided unexpected validation for RA.

*RAX Flight Part 2*—The 6-hour scenario was activated Friday morning. The scenario ran well until it was time to start up the IPS. Unfortunately, an unexpected problem occurring somewhere between the flight software and the RAX Manager caused a critical monitor value to be lost before it reached RA. The cause of this message loss has not been determined. The problem of lost monitor values could have been avoided with periodic refreshes of the monitor values. This was deemed out of scope for the purposes of the experiment, and RA was known to be vulnerable to message loss. This vulnerability led RA's estimation of the IPS state to diverge from the true state.

Fortunately, the discrepancy proved to be benign. Hence, RA was able to continue executing the rest of the scenario to achieve the rest of its validation objectives.

By executing the two flight scenarios, RAX achieved 100% of its validation objectives.

### *Effectiveness of the development and test process*


Progress in development and testing during the RAX project can be analyzed through the Problem Reports (PR) filed between April 1997 and April 1999.

Subsystem	Number of PRs
Planner/Scheduler	233
Executive	100
MIR	85
RAX Manager	22
System	77
Communication	22
Simulator	30
Others	11
Total	580

**Table 12: Number of PRs by subsystem**

A developer or a tester could file a PR, usually reporting a bug, or requesting a change in the software behavior. A few PRs were reminders of activities or checks to be performed. PRs remained open until the developers addressed them. When a resolution to the report was filed (e.g., a bug fix was provided), the originator of the report would check the validity of the resolution. If accepted the resolution was included in a formal release. A few PRs were suspended. This meant that the risk of the problem was assessed and considered acceptable within the limits of the RAX experiment.

Figure 33 gives an idea of the temporal distribution of new PRs filed over the duration of the project. The last four columns (from January 1999 to April 1999) relate to problems that were submitted to the Change Control Board process. Notice that during the number of PRs in this period is still quite high (91). This depended in part on the fact that integration with flight software started in December 1999, with RAX running on *Papabed*, and that until then RA had only been operating interacting with low fidelity simulators.



**Figure 33: Temporal distribution of Problem Reports**

PRs can be divided in three categories.

- Modeling PRs required by changes in the domain specific knowledge relative to the DS1 spacecraft subsystems;
- Engine PRs affecting the core reasoning engines of RA
- Other PRs related to other mechanisms such as the format of data file exchanged between RA components. This category also includes reminders and requests of change that were outside the scope of RAX

Figures 16, 17 and 18 describe the distribution of problems by category for each individual engine. The most stable RA subsystem was MIR. This stability manifested itself both in terms of the total number of Engine and Model PRs filed and in terms of the very few PRs of these categories filed in the last 4 months of the project. This was due both to the maturity of the MIR technology and to the fact that the problem addressed by MIR changed very little during the duration of the project.






Figure 34: Planner PRs by category




Figure 35: Executive PRs by category

The command language used by EXEC, ESL, was developed prior to the RAX project and caused a negligible number of PRs. The majority of the EXEC PRs fell into the Other category and were related to integrating the PS and MIR modules. The next largest category of PRs were model related. These tended to manifest themselves each time RA was integrated on a higher fidelity testbed. Models for EXEC were undergoing modifications quite late (February to April 1999). This was primarily due to the fact that these months covered a period of intense activity on *Papabed* with the interfaces with the details of how flight software operated being finally communicated to the RAX team. This resulted in some localized changes in interface functions and in task decomposition procedures. The effects of these changes were typically localized at the EXEC level and did not propagate up to PS models. This confirms the possibility of developing RA even on the basis of an accurate but abstract characterization of the modeled system, with much of the high-level behaviors remaining stable when further details on the behavior of the system are known.




Figure 36: MIR PRs by category

Both in the case of MIR and EXEC, testing was very effective at validating models. EXEC and MIR models have many non-interacting or loosely interacting components that can be tested independently. This reduces the number of test cases that are needed. Testing small components of the model independently like we did in RAX should scale-up to larger models for future science missions.

In the case of PS, a larger overall percentage of the PRs (about 45%) was model related. More importantly, a large number of new problems was discovered during the last four months of the project, after the formal testing process had ended. The vast majority of these problems consisted of PS operating correctly but being unable to find a plan within the allocated time limit since its search was "thrashing". These problems were particularly serious since they could easily arise in off-nominal situations during flight.

There were several reasons for this situation:

1. The ranges of some parameters turned out to be different than those assumed by PS testing, e.g., PS testing assumed turn durations were at most 20 minutes, while actual turns could take over an hour. This created stress situations not considered by formal PS testing.
2. Planning problems became more challenging when we transitioned from the 6 day scenario to the 2 day scenario. The temporal compression led to the disappearance of slack time between activities. In the 6 day scenario PS could exploit this slack to achieve subgoals without backtracking. In the 2 day scenario backtracking became necessary, revealing additional brittleness in the PS chronological backtracking search.
3. A more fundamental issue was the independence between the PS test generator and the structural



characteristics of the domain model. This led to the test generator missing a number of stress cases. For example, one problem depended upon the specific values of three continuous parameters: the time to start up the IPS engine, the time to the next optical navigation window, and the duration of the turn from the IPS attitude to the first asteroid. An equation relating these parameters can crisply characterize the stress situations. Unfortunately, automatically generated test cases based on covering pair-wise interactions of parameter values cannot reliably detect such problems.

Given the late date at which these new problems were discovered, it was not feasible to modify the test suite to test extended variations around the new baseline. Instead, we focused on just the most crucial variation: the time at which replans might occur. The objective was to ensure that the planner was robust to any re-planning contingency. Two steps were needed to accomplish this. First, the new 2 day scenario was designed to guarantee that the harmful constraint interactions of the PS domain model would be avoided under any hypothetical replanning contingency. The idea was to ensure that PS could always return a plan within the given time limit. Second, a new PS test suite was carefully designed and run to ensure that this was indeed the case.

The design methodology for this new PS test suite is instructive. Exhaustive generation of all possible plans was clearly impossible. Instead, using our knowledge of the PS model, we manually identified boundary times at which the topology of the plans would change. We identified 25 such boundary times and generated a total of 88 test cases corresponding to plans starting at, near, or between boundary times. This led to the discovery of two new bugs. Furthermore, analysis of the test results showed that PS would fail to find a plan at only about 0.5% of all possible start times. Although the probability of this failure was extremely low, contingency procedures were developed to ensure that the experiment could be successfully continued even if this PS failure actually occurred.

We used the above test suite design methodology only toward the end of RAX, after the PS model and code had been frozen. However, we believe that this (currently manual) analysis method can be generalized and extended to provide an automatic PS testing procedure throughout the development process for new application domains.

Note that the number of PRs regarding the reasoning engines of PS, EXEC and MIR was relatively small. For example, less than 10% of PS's PRs were Engine related and the last was filed in September 1998. However, the bug in EXEC encountered during the flight experiment shows that our engine validation methodology could have

improved. In fact, our testing was primarily focused on validating the knowledge in the domain models. Tests were selected to exercise the domain models. By exercising the RA on these test scenarios, we effectively tested the domain models and engines as a unit. However, especially for concurrent systems such as EXEC, a much better approach is to thoroughly formally validating the logic of the engines through the use of formal methods []. Although expensive, this form of testing can give high level of quality assurance on the core of the RA technology. Moreover, since the engines remain unchanged over a large number of applications, the cost of this testing can be amortized across several missions.


## *Costing*

Figure 37 below gives an overall view of the costing of the Remote Agent Experiment starting from October 1997 when tracking information was available. The figure describes costs based on development, testing, integration and technical management activities. And on the Y axis we show the Full Time Equivalence (FTE') exerted. Costing by FTE's is more appropriate in this case because of the differing accounting standards used at NASA ARC and JPL.

The chart clearly shows the distinct development and testing and integration efforts being partitioned in time; development efforts were clearly focused before the move to the high-fidelity testbeds. And while testing and integrations efforts were ongoing activities, they came to dominate the latter part of the move to the testbeds. While the overall trend is a curve with diminishing figures, there are some features that need some explanation.

The first peak in the Oct-Dec 97 timeframe corresponds to the time when formal test plans were put together and UNIX testing began. In addition the RAX Manager was also delivered to the flight team at this time. The peak therefore, is categorized by these efforts and the resulting testing and bug fixing that took place.

The second peak in the Aug-Oct 98 timeframe corresponds to a number of events. Primarily this was dealing with new code deliveries to the planner engine to allow the EXEC to deal more robustly with additional information in the plans. This increased effort highlights the extra individuals from outside the RA team who made these efforts possible. In addition all team members were gearing up towards testing on the papabed, the highest fidelity testbed available at that time. Subsequent to that event, the curves show a deep decline, as expected, in the development efforts when the team focused more on integration and testing on the various testbeds available. Efforts dealing with integration therefore show a perceptible increase.



**Figure 37. Remote Agent Costing**

Lastly, the gap between the testing and integration efforts appears to be inverted in the December 98 and May 99 timeframe. The primary reason for this was our late arrival on the high-fidelity testbeds. This resulted in our efforts to be redoubled with integrating on these testbeds. It was also the case that working on these testbeds took time and effort beyond what was necessary on the lower fidelity testbeds (*Unix* and *babybeds*) that were available early on. Training and detecting problems with the configuration also took up a substantial portion of our time and effort resulting in a larger manpower effort for integration as shown.

### Scaling Up

The Remote Agent experiment demonstrated all the capabilities of the RA, but the domain models encoded only what was needed for the experiment scenario. This was a small fraction of what would be needed to command all aspects of a complete science mission.

Future missions can expect that their modeling effort will exceed those of the RA experiment, though it is difficult to say what the scaling factor might be. Validation of the models also benefited from a reduced scope. We were able to focus our efforts on a single scenario and related variants. To apply the same methods to a full science mission will require either significantly more testing resources or more effective validation methods.

The time needed to generate plans will probably increase for future science missions, though that depends on plan complexity and CPU speed. The RAX plans were relatively simple, consisting of 15-25 executable activities, 50-80 tokens and 90-134 constraints. They took 50-90 minutes to generate using about 25% of the RAD6000 CPU. Generation time for more complex plans can be expected to scale roughly linearly with plan complexity (tokens and constraints). Plan complexity depends on the length of the plan and the goals that the plan is asked to achieve, both of which are under control of the operations team.

Plan generation time determines how quickly the RA can respond to failures and external events. If the re-plan duration is too long, the RA may not be able to generate a new plan in time. During normal operations the planning duration is far less important since the next plan can be generated while the current one is executing.

The remote agent reasoning engines achieved all of their validation objectives, and demonstrated all of their key capabilities. The RA as a whole demonstrated that it can command a spacecraft and respond robustly to failures. Future missions can expect that these capabilities will perform equally well on a full-scale science mission.

### Lessons Learned

The RA team learned valuable lessons in a number of areas including RA technology and processes, tools, and even autonomy benefits to missions.

**Robustness of the Basic System**—Model validation alone does not suffice; the rest of the system, including the underlying inference engines, the interfaces between the engines, and the ground tools, must all be robust. Given our resource constraints, we made the decision to focus our formal testing on model validation, with engine and interface testing happening as a side effect. This was a reasonable strategy: code that has been unchanged for years is likely to be very robust if it has been used with a variety of different models and scenarios. However, newer code does not come with the same quality assurance. Furthermore, as the deadlock bug in flight showed, subtle timing bugs can lay hidden for years before manifesting themselves.

**Conclusion:** The primary lesson is that the basic system must be thoroughly validated with a comprehensive test plan as well as formal methods, where appropriate, prior to model development and flight insertion. Interfaces between systems must be clean and well specified, with automatic code generation being used to generate actual interface code, telemetry, model interfaces, and test cases; code generation proved to be enormously helpful in those cases where we did use it.

**Robustness of Model Based Design**—As mission development times becomes shorter and mission objectives become more ambitious, it is less and less likely that an accurate model of each spacecraft component will be available early in the flight and ground software development cycle. Dealing with this uncertainty is a major problem facing future missions. By emphasizing qualitative and high-level models of behavior RA can help solve this dilemma. Qualitative, high-level models can be captured early in the mission lifetime and should need only minor adjustments when the hardware is better understood. Our experience on RAX essentially confirms this hypothesis. Initial spacecraft models used by PS, EXEC and MIR were built

early in the DS1 mission, before April 1997. During the following year and a half, EXEC and MIR models did not change and the PS model was only changed in order to support more efficient problem solving by the search engine, not in order to reflect new knowledge of the spacecraft behavior. In the last phase of the experiment preparation, when communications between the RAX team and the DS1 team resumed, adjustments were needed to finalize the interface between the low-level EXEC primitives and flight software.

**Conclusion:** Contrary to much concern, the type of qualitative, high-level models used by the RA requires little tuning throughout the project. The usefulness of the models for software development has been validated.

*Model Design, Development and Test*—One of the biggest challenges we faced was model validation. This was particularly true during validation testing, when even small changes in the models had to be carefully and laboriously analyzed and tested to ensure that there were no unexpected problems. In fact, in some cases we chose to forgo a model change, and instead decided to institute flight rules that would preclude the situation that required the model change from arising. A related issue was that methods do not yet exist to characterize RA's expected behavior in novel situations. This made it difficult to precisely specify the boundaries within which RAX was guaranteed to act correctly. While the declarative nature of RA models was certainly very helpful in ensuring the correctness of models and model changes, the difficulty stemmed from unexpected interactions between different parts of the model, e.g., different parts of the model may have been built under different, implicit, conflicting assumptions.

**Conclusion:** The central lesson we learned here was the need for better model validation tools. For example, the automated test running capability we developed proved to be enormously helpful, as it allowed us to quickly evaluate a large number of off-nominal scenarios. However, scenario generation and evaluation of test results were time consuming. In some cases, the laborious process we followed to validate model changes has provided us with concrete ideas for developing tools that would dramatically simplify certain aspects of model validation. Preliminary work in the area of formal methods for model validation is also very promising. Finally, we need to develop better methods for characterizing RA's behavior with a specific set of models, both as a way of validating those models and as a way of explaining the models to a flight team.

*Onboard Planning*—Since the beginning of RA, on-board planning has been the autonomy technology that most challenges the comfort level of mission operators. Commanding a spacecraft with high-level goals and letting it autonomously take detailed actions is very far from the traditional commanding approach with fixed-

time sequences of low-level commands. We believe that during RAX the flawless demonstration of on-board planning has provided a powerful existence proof of the feasibility of the approach. Our own discomfort with the discrepancy between tested behavior and in-flight behavior of PS during RAX was a surprising mirror of the objections of the critics of autonomy.

**Conclusion:** It is difficult to move past the mindset of expecting complete predictability from the behavior of an autonomous system. However, RAX has demonstrated that the paradigm shift is indeed possible. In the case of PS behavior during RAX, the point is not that the combination of pictures requested by NAV had never been experienced before, but that the problem-solving behavior that the planner used to achieve each individual picture goal had indeed been tested. Confidence in complex autonomous behaviors can be built up from confidence in each individual component behavior.

*Design for Testability*—System-level testing is an essential step in flight preparation. Designing the RA to simplify and streamline system-level testing and analysis can enable more extensive testing, thus improving robustness. In RAX, system-level testing proved to be cumbersome. The primary reason for this was the absence of efficient tools to generate new mission scenarios, so that all system tests had to be variations on the nominal scenarios. Hence, to test a particular variation, one was forced to run a nominal scenario up to the point of the variation, e.g., testing thruster failures during turns required at least 6 hours, since the first turn occurred about 6 hours into the scenario.

**Conclusions:** The difficulty of generating new mission scenarios is easily addressed: a graphical tool allowing visual inspection and modification of mission profiles, as well as constraint checking to ensure consistency, can dramatically simplify the construction of new mission profiles. Such a tool is now being constructed. Nonetheless, overall RA validation is still necessary to ensure that RA will properly handle each new mission profile (see below).

*Systems Engineering Tools*—Coding the domain models required substantial knowledge acquisition, which is a common bottleneck in Artificial Intelligence systems. It is better to have the domain expert code the models directly.

**Conclusion:** Develop tools and simplify the modeling languages to enable spacecraft experts to encode models themselves. Employ tools and languages already familiar to the experts. Organize the models around the domain (Attitude Control, Power, etc.) rather than around the RA technology (planner, exec, MIR ).

*Mission Profile Development*—The RA is commanded by goals specified in a mission profile. For the experiment, constructing the profile was a “black art” that only one or two people on the RA team could perform. The mission

planners and operations personnel must be able to specify goals themselves.

**Conclusion:** Simplify specification of goals. When possible, use approaches already familiar to mission planner, such as graphical timeline displays and time-ordered listings. Provide automated consistency checking.

*Adaptability to Late Model Changes*—The spacecraft requirements and operating procedures change throughout development, and even after launch. We were unable to encode late changes, due to the regression-testing overhead that each change required.

**Conclusions:** The validation cost of model-changes must be reduced. Some possibilities include tools to evaluate the consequences of model changes on testing. The models already support localized changes. Procedures are needed to uplink and install just those changes.

*Ground Tools*—Ground tools ought to be developed well in advance of the actual flight and be used as a primary means to test and understand how to operate complex systems. Given the late date of development of most of the ground tools, a good many of them felt not well integrated. As a result only the tools displaying or interpreting data in the most obvious way were of high value.

*Telemetry*—In addition to an on-board textual log file downlinked at the end of the experiment or on request, RAX sent a stream of binary telemetry packets, one for each significant event, that were displayed as color-coded text on the ground. Among other things, the telemetry allowed us to monitor all on-board communication among RAX modules and between RAX and the flight software. This proved valuable in allowing us to quickly diagnose the anomalies that occurred. We immediately knew that the reason RAX had failed to turn off the ion engine was that it had stopped executing the plan in some unanticipated manner; we knew RAX was still running and could also rule out a plan abort or a failure to send just one command. Similarly, we immediately narrowed down the second anomaly to a monitor message that was either not sent or not received.

**Conclusion:** Ensuring sufficient visibility on all platforms, including in flight, requires adequate information in telemetry. The best way to ensure this is to design the telemetry early, and to use it as the primary, if not the only, way of debugging and understanding the behavior of the system during integration, test, and operations.

*Team Structure for RA Model Development*—The RAX team was structured horizontally along engine boundaries. See Table 2. This meant that team members specialized in one of the PS, EXEC, and MIR engines, and each team was responsible for modeling all spacecraft subsystems for their engine. This horizontal organization was appropriate for RAX, since it was our first major experience in modeling spacecraft subsystems for flight.

Hence, it made sense for engine experts to do all modeling for their engine. However, this organization has several shortcomings. Perhaps the most significant shortcoming was that knowledge of any one spacecraft subsystem (e.g., attitude control, ion propulsion, MICAS camera) was distributed across the three teams; one needed discussions with three individuals to get a complete understanding of how a subsystem was commanded by RA.

**Conclusions:** These shortcomings suggest an alternate structuring for a future SW team. Instead of a horizontal structure, teams might be organized vertically along spacecraft subsystem or domain unit boundaries; e.g., a single team would be responsible for developing all models for the attitude control system. This would ensure internal coherence of the resulting model. Furthermore, since modelers would need to understand how to use all three engines, they can make well motivated decisions on how best to model a subsystem to exploit the strengths of each engine and avoid information duplication.

While a vertical team organization has its benefits, certain aspects of model development intrinsically involve managing and reasoning about global constraints, e.g., power allocation strategies, system-level fault protection. Hence, it is important to involve systems engineers to develop these global strategies.

## FUTURE APPLICATIONS

Future work regarding Remote Agent can be divided into three categories: fundamental improvements in the capabilities of its components, improvements in usability or deployability, and upcoming demonstrations or applications. Since the experiment a significant effort has gone into basic research to improve future iterations of Remote Agent. For example, more capable version of Livingstone has been developed that better handles ambiguity when tracking the state of the spacecraft. Livingstone now tracks a number of most likely states the spacecraft could be in, given the observations it has received thus far. If new observations invalidate the possible states MIR considered most likely, it re-analyzes the commands that have been given and the possible failures in order to determine which previously unlikely states now explain the unexpected observations.

PS has a number of efforts underway to improve the underlying software implementation - it now has a new modular software architecture which allows plugging in of various search techniques in the engine, work is underway in model analysis which will allow early detection of domain model inconsistencies, analysis of static models is also being undertaken to automatically generate search control instrumentation. The latter approaches will allow rapid prototype development of planner models by non-technologists using incremental

model development via "what-if" analysis to vastly reduce development costs. It will also provide mission staff with a better understanding of how autonomy architectures will fit into the overall design of the FSW.

Other efforts are also in place to redesign the system architecture to allow the EXEC access to the planner temporal database and algorithms. A unified modeling language is being developed with cleaner semantics to allow EXEC to respond to exogenous events more rapidly.

Applying Remote Agent to the Deep Space 1 spacecraft provided a wealth of practical lessons about what was needed to create a sustainable autonomy engineering process and make this technology usable for main-line mission development and operations. PS and MIR have been re-architected, modularized and implemented in C++ rather than Lisp. These next-generation versions are in alpha testing at the date of this report. EXEC is expected to be re-architected and implemented in C by the end of calendar year 2000. The Remote Agent team is now developing tools for graphically creating and debugging models, for automating much of the integration of Remote Agent with traditional flight software, and for allowing humans and autonomous software to interact more easily. The team is collaborating with software verification researchers at NASA Ames Research Center and Carnegie-Mellon University to allow certain Remote Agent models to be analyzed to prove they cannot recommend undesired behavior. In short these research and development efforts are designed to make the Remote Agent and similar technologies more capable, easier to use, and easier to test and validate.

Remote Agent technology is successfully being transferred beyond the original team and several groups are currently building prototypes with Remote Agent in order to evaluate it. At NASA's Kennedy Space Center, Remote Agent applications are being developed to evaluate RA for missions involving in-situ propellant production on Mars on the 2003 lander or a future piloted mission. Applications for shuttle operations are being pursued as well. At the Jet Propulsion Laboratory, Remote Agent is being evaluated as the baseline autonomy architecture for the Origins Program Interferometry Instruments and is being used in the JPL interferometry testbed. The New Millennium Program's Deep Space Three, a space-based interferometry mission which includes two or three spacecraft cooperating to take science observations, may be one early customer of this development. At Johnson Space Center, components of Remote Agent are being integrated into an ecological life support testbed for human missions beyond Earth orbit. At Ames Research Center, Remote Agent technology is being incorporated into software for more robustly controlling planetary rovers. Working with Orbital Sciences Corporation, Ames is working to demonstrate Remote Agent as it applies to streamlining the checkout and operation of a reusable launch vehicle. This

demonstration will fly on the X-34 vehicle. A similar experiment will be flown on the X-37 vehicle in collaboration with Boeing.

## ACKNOWLEDGEMENTS

This report describes work performed at the NASA Ames Research Center and at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration.

The Remote Agent Experiment would not have been possible without the efforts of the DS1 flight team and Harlequin Inc.. In addition we gratefully acknowledge the direct contribution of the following individuals:

Steve Chien, Micah Clark, Scott Davis, Julia Dunphy, Chuck Fry, Erann Gat Ari Jonsson, Ron Keesing, Guy Man, Sunil Mohan, Paul Morris, Barney Pell, Chris Plaunt, Greg Rabideau, Scott Sawyer, Reid Simmons, Mike Wagner, Brian Williams, Greg Whelan and David Yan

In addition, the following individuals were instrumental in providing valuable support and encouragement throughout the duration of the development and flight of the Remote Agent:

Abdullah Aljabri, Martha Del Alto, Ralph Basilio, Magdy Bareh, Kane Casani, Rich Doyle, Dan Dvorak, Dan Eldred, Julio Fernandez, Peter Gluck, Jack Hansen, Ricardo Hassan, Lauraine Fesq, Ken Ford, Sandy Krasner, David Lehman, Frank Leang, Mike Lowry, Nicole Masjedizadeh, Maurine Miller, Alex Moncada, Mel Montemerlo, Peter Norvig, Keyur Patel, Bob Rasmussen, Marc Rayman, Mark Shirley, Martin Simmons, Helen Stewart, Gregg Swietek, Hans Thomas, Phil Varghese and Udo Wehmeier.

We also gratefully acknowledge the participation and support of the following organisations: Caelum Research Corp, and Recom Technologies

## REFERENCES

- [1] Erann Gat, "ESL: A language for supporting robust plan execution in embedded autonomous agents," Proceedings of the AAAI Fall Symposium on Plan Execution, AAAI Press, 1996.
- [2] Erann Gat and Barney Pell, "Abstract Resource Management in an Unconstrained Plan Execution System," in Proc. of IEEE Aeronautics (AERO-98), Aspen, CO, IEEE Press, 1998.
- [3] Barney Pell, Erann Gat, Ron Keesing, Nicola Muscettola, and Ben Smith, "Robust Periodic Planning and Execution for Autonomous Spacecraft," In Procs. of IJCAI-97, 1997.
- [4] Brian C. Williams and P. Pandurang Nayak, "A model-based approach to reactive self-configuring systems." In Proceedings of AAAI-96, pages 971-978, 1996.

- [5] Barney Pell, Ed Gamble, Erann Gat, Ron Keesing, James Kurien, Bill Millar, P. Pandurang Nayak, Christian Plaunt, and Brian Williams, "A hybrid procedural/deductive executive for autonomous spacecraft," In P. Pandurang Nayak and B. C. Williams, editors, *Procs. of the AAAI Fall Symposium on Model-Directed Autonomous Systems*, AAAI Press, 1997.
- [6] N. Muscettola, "HSTS: Integrating planning and scheduling," in Fox, M., and Zweben, M., eds, *Intelligent Scheduling*, Morgan Kaufman.
- [7] N. Muscettola, B. Smith, S. Chien, C. Fry, G. Rabideau, K. Rajan, D. Yan, "On-board Planning for Autonomous Spacecraft," in *Proceedings of the fourth International Symposium on Artificial Intelligence, Robotics and Automation for Space (i-SAIRAS 97)*, July 1997.
- [8] Muscettola N., P. Nayak., B. Pell, and B. Williams, "Remote Agent: To Boldly Go Where No AI System Has Gone Before," *Artificial Intelligence* 103(1-2):5-48, August 1998
- [9] J. de Kleer and B. C. Williams, "Diagnosing Multiple Faults," *Artificial Intelligence*, Vol 32, Number 1, 1987.
- [10] D. S. Weld and J. de Kleer, *Readings in Qualitative Reasoning About Physical Systems*, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1990.
- [11] A. Jonsson, P. Morris, N. Muscettola, K. Rajan and B. Smith "Planning in Interplanetary Space: Theory and Practice", to appear in the 5<sup>th</sup> International AI Planning Systems, AIPS2000, Breckenridge, CO, 2000.
- [12] B. Smith, K. Rajan and N. Muscettola "Knowledge Acquisition for the Onboard Planner of an Autonomous Spacecraft", *European Knowledge Acquisition Workshop*, Spain 1997.
- [13] Remote Agent Web Page at <http://rax.arc.nasa.gov>.
- [14] K. Rajan, M. Shirley, W. Taylor and B. Kanefsky, "Ground Tools for Autonomy in the 21<sup>st</sup> Century", to appear in *IEEE Aerospace*, Big Sky, MT, March 2000.
- [15] Bernard D., et al., "Design of the remote agent experiment for spacecraft autonomy." In *Proceedings of the IEEE Aerospace Conference*, March 1998.
- [16] Bernard, D. et al, "Spacecraft Autonomy Flight Experience: The DS1 Remote Agent Experiment", *AIAA'99*, Albuquerque, NM, October 1999.
- [17] Nayak P. et al, "Validating the DS1 Remote Agent Experiment," *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS-99)*, Noordwijk, Holland, 1999.



## Appendix A: Telemetry Channels

The bulk of RAX monitoring and validation during the experiment was from the RAX telemetry on APID 9 & 10, channels W-500 to W-570, and the downlinked log files.

Channel	Mnemonic
W-500 - W-570	(RAX channels)
P-0300	LPE_PASM_mgr
APID 9 & 10	Monitored RAX behavior. Packets were in a RAX-specific format.
APID 45	Log files downlinked after the experiment (plan files and detailed execution trace).

The following channels were also activated for RAX:

Channel	Mnemonic
F-1048	FaultEnaStat
F-1052	BusSCstatus
F-1055	IPS_SCstatus
F-1057	PDS_SCstatus
F-1058	ACS_SCstatus
F-1060	RAX_SCstatus
F-1063	BusGDstatus
F-1066	IPS_GDstatus
F-1068	PDU_GDstatus
F-1069	ACS_GDstatus
F-1071	RAX_GDstatus
D-0149	buf_pkt_09
D-0150	sent_pkt_09
D-0165	buf_pkt_10
D-0166	sent_pkt_10
F-0716 - F-0727	

## Appendix B: DS1 Technology Validation Power On Times

- The Remote Agent Experiment first ran from May 17, 1999, 5am PST to Wed 5/18/99, 7pm PST.
- It ran again from May 21, 1999, 7:15am PST to 1:30 pm PST (RAX\_STOP).
- The log files were downlinked by May 21, 1999 4:00 PST.

## Appendix C: Acronym Definitions

AutoNAV .....	Autonomous Navigation subsystem of FSW
ACS.....	Attitude Control Subsystem of FSW
APE.....	Attitude Planning Expert subsystem of FSW
ARC .....	Ames Research Center
CCB .....	Change Control Board
CPU.....	Central Processing Unit (computer)
DDL .....	PS Domain Description Language
DMD.....	**** not defined in paper
DS1 .....	Deep Space One spacecraft



DSN .....Deep Space Network  
 ESL .....Executive Support Language  
 EXEC.....Remote Agent Smart Executive  
 FTE .....Full Time Equivalent  
 FSW .....DS1 Flight Software  
 GMT .....Greenwich Mean Time  
 HGA.....High Gain Antenna \*\*\* not defined in paper  
 HSTS.....\*\*\*\*\*not defined in paper  
 HV.....High Voltage \*\*\*\*\*not defined in paper  
 IPCO\_RAXO.....\*\*\*\* not defined in paper  
 IPS.....Ion Propulsion System  
 IRS .....Incremental Refinement Scheduler  
 JPL .....Jet Propulsion Laboratory  
 LGA .....Low Gain Antenna \*\*\* not defined in paper  
 MGDS.....\*\*\* not defined in paper  
 MICAS.....Miniature Integrated Camera And Spectrometer  
 MI.....Mode Identification component of MIR  
 MIR.....Remote Agent Mode Identification and Recovery module (Livingstone)  
 MR .....Mode Recovery component of MIR  
 MM .....Remote Agent Mission Manager module  
 NASA .....National Aeronautics and Space Administration  
 NewMAAP.....New Millennium Autonomy Architecture rapid Prototype  
 OD.....Orbit Determination  
 OPNAV .....Optical Navigation Module subsystem FSW  
 PASM .....\*\*\* not defined in paper  
 PEF .....Predicted Events File  
 PEPE.....\*\*\*\*\* not defined in paper  
 PR.....Problem Report  
 PS .....Remote Agent Planner/Scheduler  
 RA.....Remote Agent  
 RAX.....Remote Agent Experiment  
 RAXM.....RAX Manager  
 RT .....Remote Terminal  
 SSA .....\*\*\* not defined in paper  
 TDB .....HSTS Temporal Database  
 TCV .....Thrust Vector Control \*\*\* not defined in paper  
 UV.....Under Voltage \*\*\*\*\* not defined in paper

---